FY2005 Accomplishments

Applied Mathematics

Improving the Performance of Mesh Optimization Techniques

Paul Hovland and Todd Munson*, Argonne National Laboratory

Summary

Optimizing the element quality in a mesh can significantly reduce the time required to find an approximate solution to a partial differential equation. Calculating the improved mesh in the least amount of time requires both an efficient algorithm and an efficient implementation. The code Argonne researchers have developed to compute an optimal mesh has successfully reduced the execution time of a spectral-element method applied to a fluid dynamics application by 30 percent.

Discretization methods are common techniques for computing approximate solutions to partial differential equations. These methods decompose the given domain into a finite set of elements, triangles, or tetrahedrons, for example, to produce a mesh used within the approximation scheme. Several factors affect the accuracy of the solution to the partial differential equation computed: the degree of the approximation scheme and the number of elements in the mesh, and the quality of the mesh. Optimizing the quality of the mesh prior to computing the approximate solution can improve the condition number of the linear systems solved, reduce the time taken to compute the solution, and increase the numerical accuracy.

The savings in computational time from using the optimized mesh can be substantial. One application we investigated with Paul Fischer (ANL) and S. Balachandar and Lin Zhang (UIUC) was to solve the Navier-Stokes equations for a fluid containing a dilute suspension of particles. The approximate solution was obtained by applying a spectral-element method to a hexahedral mesh. The top of Figure 1 depicts their original mesh, while the bottom shows the mesh after shape-quality improvement techniques have been applied. The original mesh has many regular elements, while the optimized mesh loses much of this structure. However, their spectral-element method required 29 hours to compute a solution when using the original mesh, but only 20 hours when using the optimized mesh, a 30 percent reduction in time.

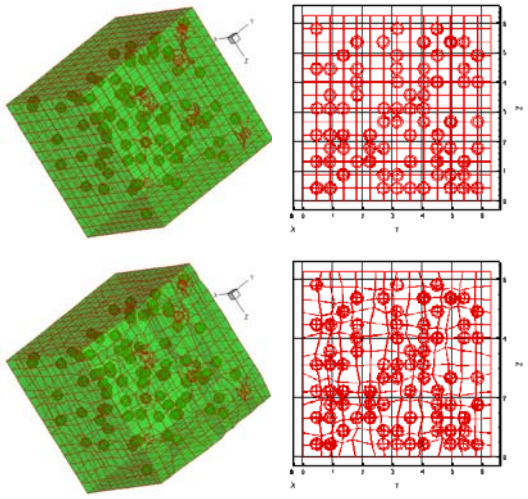

Figure 1. Original mesh and side view (top), and optimized mesh and side view (bottom) for a fluid dynamics application with a dilute suspension of particles.

The optimization problem we solve computes positions for the vertices in the

* Mathematics and Computer Science Division, (630) 252-4279, tmunson@mcs.anl.gov

mesh to improve the average element quality by using the inverse mean-ratio metric, a shape-quality metric measuring the distance between a trial element and an ideal element, a regular tetrahedron, for example. The objective function for the resulting optimization problem is nonconvex and consists of the sum of many fractional terms. Included in the optimization problem are periodic boundary conditions and constraints restricting the vertices to the planes and spheres defining the domain.

The computational properties of an inexact Newton method with a line search have been extensively studied when solving a simplified optimization problem where the positions of the vertices on the boundary of the mesh are fixed. The conjugate gradient method with a block Jacobi preconditioner is applied to solve the systems of equations. To improve the performance of this code, we have applied several techniques: using a block sparse matrix data structure to store the Hessian, reordering the problem data, and preconditioning the iterative method. These techniques can significantly reduce the computational time, especially for large problem instances.

Modern microprocessors are highly sensitive to the spatial and temporal locality of data sets. Therefore, reordering the vertices and elements in a mesh can have a significant impact on performance. We have developed metrics and models for mesh reordering and have investigated the performance of several reordering algorithms. Modeling the mesh as a hypergraph is critical and can lead to performance improvements of nearly 50 percent. Figures 2 and 3 show the effects of the data-reordering algorithms on the sparsity pattern and execution time. The performance of the gradient and Hessian evaluations was improved by approximately 20 percent using the reverse mode of automatic differentiation. Another 10 percent gain in overall performance was achieved with an improved implementation of the reciprocal cube root function.

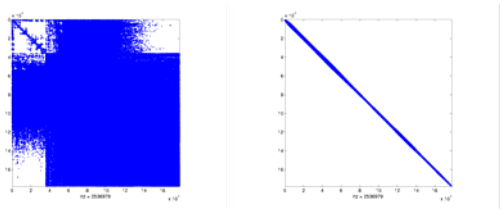

Figure 2. Sparsity pattern of the Hessian matrix for original (left) and reordered problems (right) for the duct8 mesh.

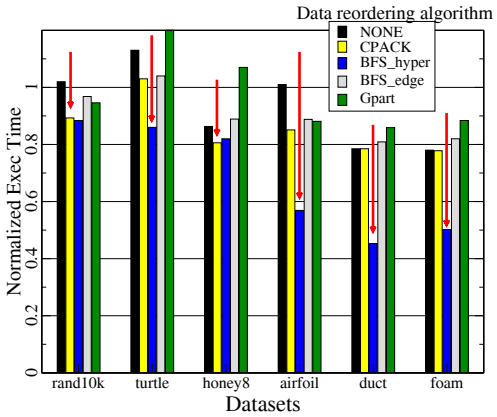

Figure 3. Effects of data-reordering algorithms on optimization time.

We plan to extend our work to constrained optimization problems where certain vertices are required to lie on the boundary of the domain. This work will require either access to the geometry or a mechanism to infer geometric features from the mesh data. These constrained problems are more difficult to solve. We will investigate techniques to improve performance of the resulting code.

For further in formation on this subject contact:

Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

A Family of Mimetic Finite Difference Methods on Polygonal and Polyhedral Meshes

Franco Brezzi[†], University of Pavia,
Konstantin Lipnikov[*], Los Alamos National Laboratory,
Valeria Simoncini[‡], University of Bologna

Summary

We gave a rigorous mathematical description of a rich family of mimetic finite difference discretizations for diffusion problems on polygonal and polyhedral meshes. We developed a computationally cheap method for generating particular members of this family.

In many applications, the mathematical model is formulated initially as a system of first-order partial differential equations, with each equation having a natural connection to physical aspects of the problem. For the diffusion problem these equations are

$$\text{div } F = b, \qquad F = -K \text{ grad } p,$$

which describe the mass conservation and the Darcy law, respectively. The unknown variables are pressure p and flux F. The material properties are described by a full symmetric tensor K.

There are many discretization schemes with equivalent properties which can be used to solve the diffusion problem. The mimetic finite difference (MFD) method is well suited for solving this first-order system, since it preserves essential properties (symmetry and mass conservation) of the continuum model. In article [1], we employed an innovative technique to give the first rigorous mathematical description of a rich family of MFD methods. This family will allow us to tackle other computational problems such as enforcement of the discrete maximum principle.

We also developed a novel computationally inexpensive algorithm for deriving particular members of the family of MFD methods. These members are described by a single parameter. With this algorithm, solving the diffusion problems of a polyhedral mesh is as simple as on a tetrahedral mesh.

* Theoretical Division, Group T-7, MS B284, lipnikov@lanl.gov; [†]Department of Mathematics, brezzi@imati.cnr.it; [‡]Department of Mathematics, valeria@dm.unibo.it

The illustrative example below shows that there is a big interval for the parameter u where the discretization errors vary only 3 times.

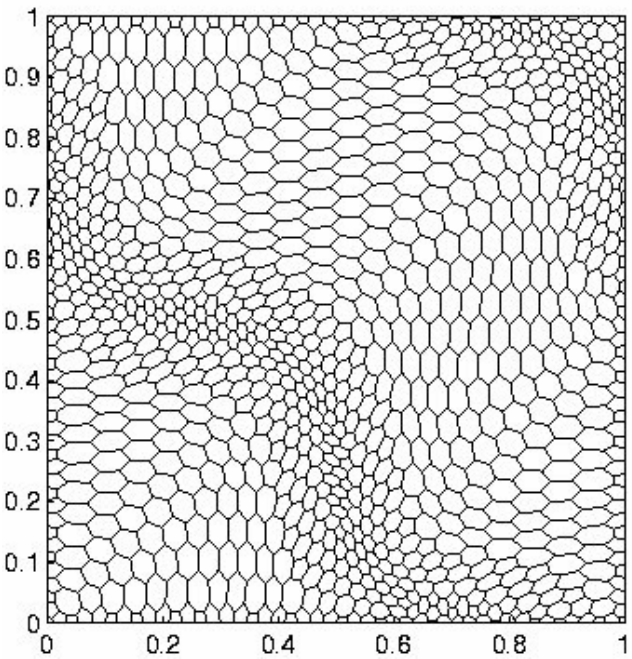

The polygonal mesh used in the convergence study.

[1] Brezzi, F., K. Lipnikov, M. Simoncini, *Math. Mod. Meth. Appl. Sci.* (2005) **15,** pp. 1533-1552.

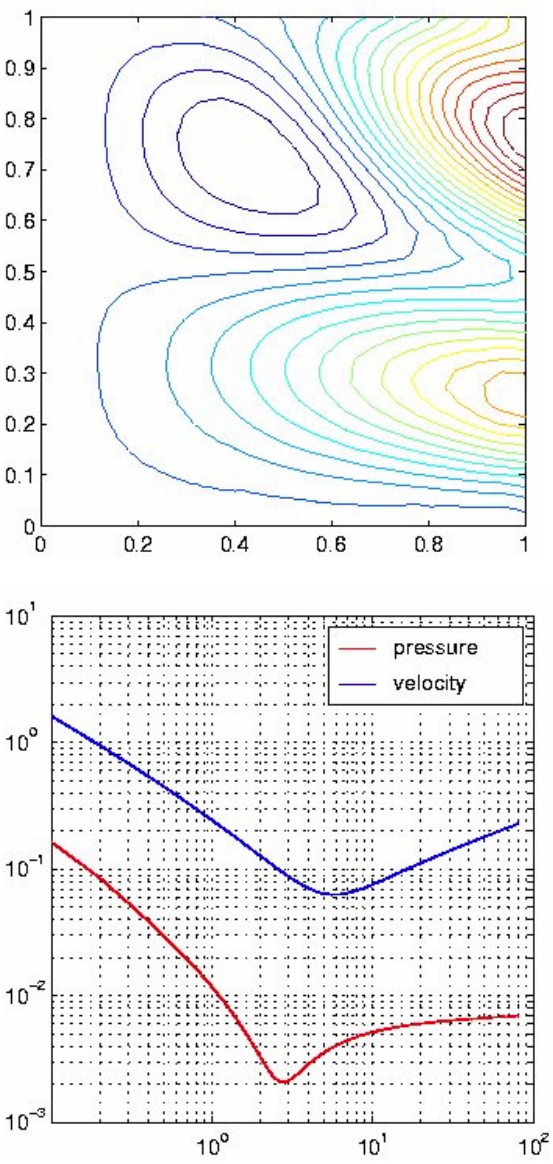

The top picture shows the solution isolines. The bottom picture shows the mesh dependent L_2-norm of errors for pressure p and flux F (vertical axis) as functions of parameter u (horizontal axis).

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

"Mesh Quality Improvement & Optimization"

Patrick M. Knupp*,
Sandia National Laboratories

Summary

The goal of this project is to provide basic mathematical research on mesh optimization methods for improvement of simulation accuracy, efficiency, and stability. A new paradigm for mesh optimization has been developed which, for the first time, encompasses the many varieties of mesh optimization goals within a single theoretical framework. The framework provides fresh insight into mesh improvement algorithms and their theoretical foundation. The paradigm employs advanced metrics and target-matrices to provide a crisp conceptual mapping between application goals and mesh optimization algorithms. Portions of the new paradigm are implemented in the Mesquite code and have been successfully applied to deforming mesh problems at SLAC, SNL, and CSAR.

Suitable computational meshes can favorably impact simulation accuracy, efficiency, and stability. Suitable meshes can be obtained in a variety of ways, including mesh smoothing and optimization. Optimization provides a rigorous mathematical framework for the improvement of meshes that is lacking in other approaches. In mesh optimization one formulates a local mesh metric which measures the 'quality' of the mesh. Local metrics over the mesh are assembled into a mesh-wide objective function which can be optimized to improve mesh vertex locations or local topology. A great deal of progress on the formulation of local quality metrics that measure quality relative to the application through the use of target-matrices has been made recently by researchers at Sandia National Labs.

Prior to the present work, mesh 'quality' for unstructured meshes has primarily focused on mesh generation issues lying outside the context of solution-adaptive meshing. Thus, for example, quality metrics were based on assumptions about ideal element shapes. The resulting meshes were essentially created independently of specific applications, leaving room for considerable improvement.

A better approach to mesh improvement is to enable application experts to define mesh quality in terms of their particular simulation accuracy and efficiency requirements. The target-matrix paradigm addresses this need though the use of 'target matrices' which define local mesh quality in relative terms meaningful to the application instead of the absolute terms used in *a priori* meshing. A variety of application-dependent target construction algorithms have been devised that incorporate application-specific information to provide more effective mesh control. For example, in simulations involving deforming domains, targets can be constructed from the mesh on the un-deformed domain. The resulting mesh on

* 505-284-4565, pknupp@sandia.gov

the deformed domain retains the features of the un-deformed mesh as far as possible. The available information is not always enough to fully determine targets; the paradigm compensates for this with new metrics which are invariant to key properties such as element shape, size, and orientation. The paradigm is mathematically rich and has required extensive analysis to answer important questions concerning convexity, symmetries, and completeness.

The paradigm naturally addresses many application goals of optimization including deforming meshes, ALE rezoning, r-type adaptivity, mesh fix-up, and anisotropic mesh improvement. Considerable elaboration and development of the paradigm will be required in the coming years to fully demonstrate and realize its potential. In the short-term, however, successful applications of the paradigm have already occurred via their implementation in the Mesquite (Mesh Quality Improvement Toolkit) software library. Successful applications include (i) incorporation of the deforming mesh algorithm within the accelerator shape optimization capability under development at SLAC (Figure 1), (ii) incorporation of the Mesquite rezone algorithm into the Alegra code at SNL, and (iii) use of the deforming mesh algorithm at the UIUC CSAR Rocket Center for propellant burn (Figure 2).

International Linear Collider

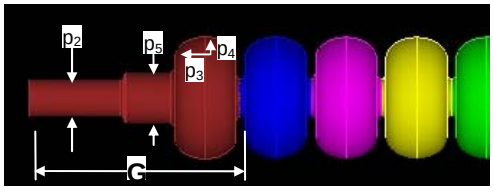

Fig. 1. Shape Optimization (SLAC)

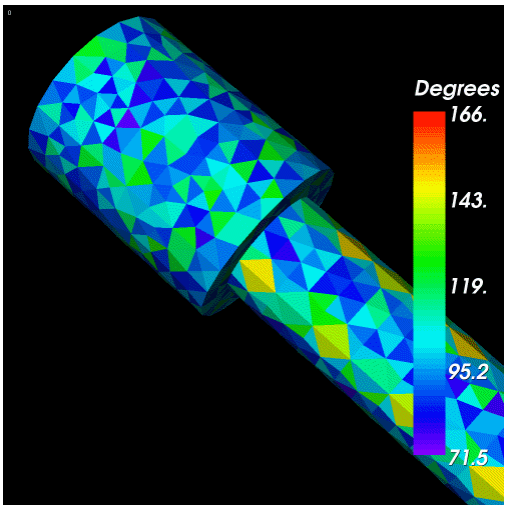

Fig. 2. Rocket Propellant Burn Simulation (CSAR)

The deforming mesh algorithm in Mesquite is used to provide gradual mesh movement for updating the mesh to changes in the accelerator cavity design parameters. The algorithm makes extensive use of target matrices to ensure the resulting mesh is valid (i.e., untangled) and bears a close resemblance to the initial good-quality mesh on the un-deformed geometry. This constitutes an advance over the commonly used elastic-springs model of mesh deformation. The deforming mesh algorithm is stressed more severely in the propellant burn simulation because significantly larger domain deformations are possible. The method is helping CSAR analysts to maintain good quality mesh during propellant burn while preserving mesh topology. By preserving topology, simpler algorithms can be used by the analysis code.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
gary.johnson@science.doe.gov

Fast Adaptive Algorithms for Fluid Particle Systems

William D. Henshaw*
Lawrence Livermore National Laboratory

Summary

There are a wide variety of problems that involve the motion of solid particles immersed in a fluid or gas. These include fluidized beds used in many industrial applications and pollution control devices that separate contaminants from the air. The modeling and simulation of these systems is important to better understand their behavior and to improve their design. New fast adaptive numerical algorithms have been developed to computationally model fluid particle systems.

As part of the Applied Math Research program effort, researchers at LLNL have developed fast adaptive algorithms for modeling the interactions between fluids and embedded solid particles. The accurate simulation of moving particles immersed in a fluid or gas is a computationally challenging task. The efficient new algorithms are based on a domain decomposition approach that combines a locally optimal representation of the particles with an efficient representation of the dynamically evolving region occupied by the fluid. The use of algorithms that adapt to the local behavior of the flow field allow limited computational resources to be dynamically targeted to the critical regions. The approach has been used to simulate experiments of gas phase combustion and condensed phase solid explosives. The resulting technology also has great potential for treating a wide class of difficult industrial problems such as the direct simulation of particulate flows, modeling of fluidized beds, and the modeling of pollution control devices.

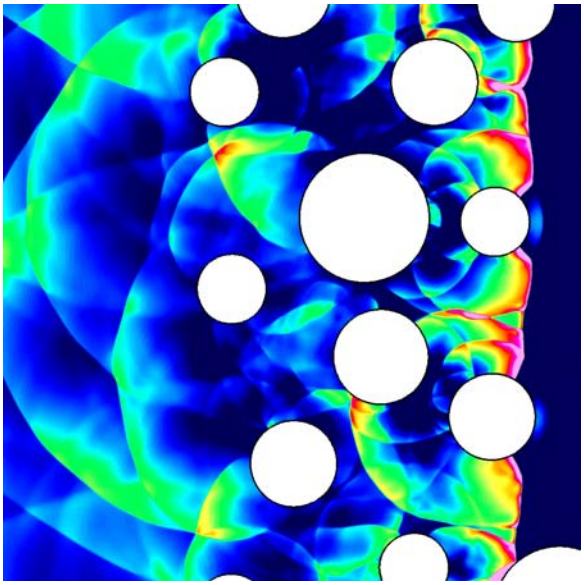

Figure 1: Particles accelerated by a compressive detonation wave that moves from left to right. The gas pressure is shown.

There are various mathematical equations that govern the motion of particles immersed in a fluid or gas. In the present case the fluid is modeled using the reactive Euler equations; this system of nonlinear hyperbolic equations embodies the laws of

* Center for Applied Scientific Computing, (925) 423-2697, henshaw1@llnl.gov

conservation of mass, momentum and energy. The translation and rotation of the particles is governed by the Newton-Euler equations of motion. The particles interact with the fluid through surfaces forces. Pressure gradients in the fluid will accelerate the particles. These mathematical equations are solved using sophisticated algorithms and numerical approximations. Careful mathematical analysis is required to show that the approximations are accurate and stable.

The domain occupied by the fluid is discretized with a Cartesian grid. In general, a grid defines the points in space where the mass, velocity, temperature and chemical properties of the fluid are represented. The finer the grid, the better the approximation is. Fine grids, however, require greater computational resources, and thus there is a trade-off between the accuracy of the answer and the time it takes to compute a solution. Each solid particle is represented with a local boundary fitted grid. The grids for the fluid and the grids around the particles overlap and are coupled through interpolation conditions. As the particles move in the fluid the grids associated with each particle evolve to follow them. With many methods it can be computationally expensive to move the grid points. By moving only the local grids associated with each particle, this expense is dramatically reduced.

If the fluid is accelerated by the motion of a boundary or particle, or if chemical reactions are occurring, then shocks and detonations may form. These features represent regions where there are extremely rapid changes in the fluid properties. A detonation in a gas, for example, may be a fraction of a millimeter wide and travel at speeds of 1500 meters per second. New grids are added (adaptive mesh refinement) in order to better resolve shocks and detonations. These new grid points are only added in regions where they are needed, thus improving the accuracy of the numerical solution without wasting computational resources. The numerical methods are implemented in computer programs that use C++ at the high-level and Fortran for low-level optimized kernels. The software is built upon the Overture object-oriented framework. The current application combines many different components of the framework including the support for moving grids, adaptive mesh refinement and overlapping grid generation.

Figure 1 shows results from the computation of a detonation traveling through a gas that contains a collection of solid particles. The particles are accelerated by the detonation wave. The particles move and compress the gas ahead. This compression leads to the creation of a small bow shock which is quickly followed by the formation of a local detonation ahead of each particle.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

High-Order Embedded Boundary Methods for Wave Propagation

Anders Petersson* and Bjorn Sjogreen
Lawrence Livermore National Laboratory

Summary

The primary goal of this project is to advance computational science by developing and analyzing numerical methods for wave propagation simulation in complex geometries. We use the second-order formulation which is approximated by finite differences on Cartesian grids. Complex geometrical boundaries are handled using the embedded boundary technique. This approach offers many advantages over the standard first-order formulations on curvilinear grids. Our research encompasses both fundamental analysis of the stability and accuracy of the new method as well as practical application to problems of interest to DOE.

Wave propagation phenomena are important in many DOE applications such as forecasting ground motions and damages due to seismic activities, non-destructive testing, underground facilities detection, oil and gas exploration, predicting the electro-magnetic fields in accelerators, and acoustic noise generation. There are also future applications that would benefit from simulating wave propagation, such as carbon sequestration via seismic reflection and geothermal energy applications.

Traditionally, linear wave propagation has been simulated using finite difference (FD) methods on one Cartesian grid that covers the entire computational domain. This technique is highly efficient both in terms of computational effort and memory requirements, and it also produces accurate solutions for box-shaped geometries and smoothly varying wave speeds. While the traditional FD method can be modified for problems with complex geometries whose boundaries or internal layers are not aligned with the grid directions, it often has difficulties with stability and accuracy. To overcome these weaknesses, we are developing improvements of the basic FD technique that allow a fully second-order accurate treatment of boundary and jump conditions in complex domains. In our approach, the partial differential equations are approximated by finite differences on Cartesian grids in second-order formulation, and the complex geometry is handled by the embedded boundary technique.

Centered finite differences are used to discretize the wave equation both in space and time. Of particular importance are details of how the boundary conditions are discretized because they determine the accuracy and stability of the numerical method. In the current approach, we discretize boundary conditions such that both the solution and its gradient are second-order accurate. Furthermore, both Dirichlet and Neumann boundary conditions are discretized in a way which avoids stiffness

* Center for Applied Scientific Computing, (925) 4-3804, andersp@llnl.gov

due to small cells cut by the boundary, without sacrificing the second-order accuracy of the method. As a result we can use an explicit time-integration method where the time-step essentially equals that of a periodic domain.

Our recent research has focused on extending our basic approach to two-dimensional Maxwell's equations. We have developed a second-order accurate scheme for domains external to perfectly conducting objects of complex shape. The boundary conditions are discretized using the aforementioned embedded boundary technique and the computational domain is truncated by a perfectly matched layer (PML). We derived estimates for the errors due to the reflections at the outer boundary of the PML and due to discretizing the continuous PML equations. Using these estimates, we showed how the parameters of the PML can be chosen so that the discrete PML solution converges to the solution of Maxwell's equations on the unbounded domain as the grid size goes to zero. This technique was used to perform scattering analysis around two-dimensional objects of complex shape (see Figure 1).

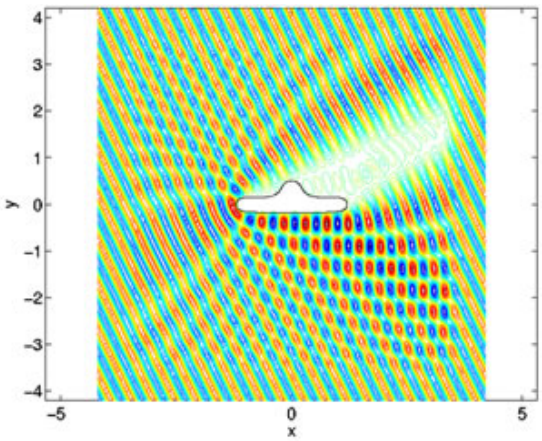

In addition, we have generalized our techniques for the scalar wave equation to handle a discontinuous jump in wave speed across an internal boundary or interface. A theory was developed for proving stability of the one-dimensional method, and numerical examples were given for the two-dimensional case where the interface is embedded into the mesh. In particular, the method was used to study electro-magnetic scattering of a plane wave by a dielectric cylinder and the numerical solutions were evaluated against the century-old analytical solution due to Mie. Point-wise second-order accuracy was confirmed.

Figure 1: The electric field resulting from an incoming planar wave scattered by an unidentified object with a perfectly conducting boundary. The computational domain is truncated by a PML where the scattered field is damped.

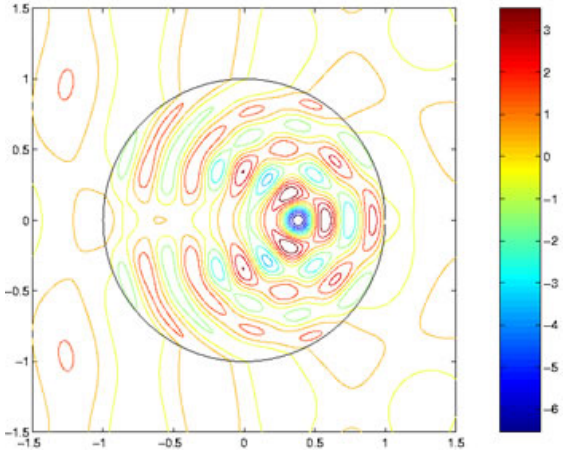

Figure 2: A planar electro-magnetic wave scattered and transmitted into a dielectric cylinder where the wave speed is Sqrt(10) smaller than the ambient media. Notice the focusing of the wave inside the cylinder.

For more information on this project, see http://www.llnl.gov/CASC/serpentine/.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

Differential Equations in Random Domains

Daniel M. Tartakovsky*, Los Alamos National Laboratory
Dongbin Xiu, Purdue University

Summary

Physical phenomena in domains with rough boundaries play an important role in a variety of applications ranging from surface imaging to manufacturing of nanodevices. Often the topology of such boundaries cannot be accurately described in all of its relevant details due to either insufficient data, or measurement errors, or both. In such cases, this topological uncertainty can be efficiently handled by treating rough boundaries as random fields, so that an underlying physical phenomenon is described by deterministic or stochastic differential equations in random domains. To deal with this class of problems, we developed a novel computational framework, which is based on stochastic mappings to transform the original deterministic/stochastic problem in a random domain into a stochastic problem in a deterministic domain. The latter problem has been studied more extensively, and existing analytical/numerical techniques can be readily applied.

Given a proper spatial resolution, virtually any natural or manufactured surface becomes rough. Consequently, there is a growing interest in experimental, theoretical and numerical studies of deterministic and probabilistic descriptions of such surfaces and of solutions of differential equations defined on the resulting domains.

The early attempts to represent surface roughness and to study its effects on the system behavior were based on simplified, easily parameterizable, deterministic surface inhomogeneities, such as symmetrical asperities to represent indentations and semispheres to represent protrusions. Alternatively, one can use random fields to represent rough surfaces whose detailed topology cannot be ascertained due to the lack of sufficient information and/or measurement errors. We adopted random representations of rough surfaces, because of their generality. Such an approach allows one not only to make predictions of the system behavior, but also to quantify the corresponding predictive uncertainties.

The presence of uncertainty in rough boundaries necessitates the development of new approaches for the analysis and numerical solution of differential equations defined on random domains. For example, it has been demonstrated that classical variational formulations might not be suitable for such problems, and finite difference approaches remain accurate only for relatively simple rectangular irregularities.

The adoption of a probabilistic framework to describe rough surfaces makes even an essentially deterministic problem stochastic,

* Theoretical Division, (505) 667-0968, dmt@lanl.gov; also at University of California, San Diego

e.g., deterministic equations in random domains give rise to stochastic boundary-value problems. This necessitates the search for new stochastic analyses and algorithms. For example, one of the very few existing numerical studies has employed traditional Monte Carlo simulations, which have turned out to be so computationally expensive, as to become impractical.

Recently, we presented a computational framework that is applicable to a wide class of deterministic and stochastic differential equations defined on domains with random (rough) boundaries. A key component of this framework is the use of robust stochastic mappings to transform an original deterministic or stochastic differential equation defined on a random domain into a stochastic differential equation defined on a deterministic domain. This allowed us to employ the well-developed theoretical and numerical techniques for solving stochastic differential equations in deterministic domains.

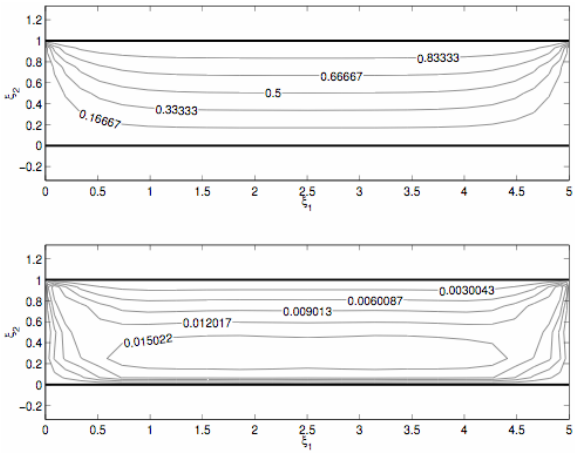

Figure 1: Predicted concentration distribution (a) and the corresponding predictive error bounds (b).

In [1], we analyzed diffusion in the rectangle with the rough random bottom. The geometric uncertainty translates into the predictive uncertainty. Hence the best estimate of the concentration of a diffusing substance must be accompanied by a measure of the corresponding predictive uncertainty. This is accomplished by computing the mean and standard deviation of concentration, both of which are shown in Figure 1.

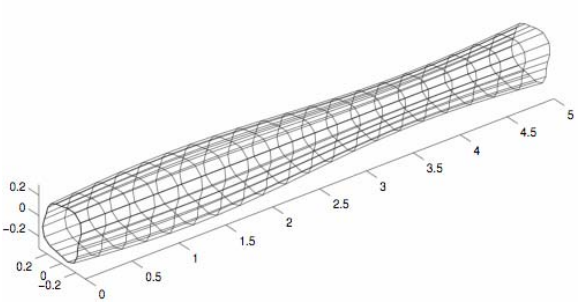

Figure 2: A schematic representation of a tube with a rough surface.

In [2], we conducted a similar analysis of transport in Stockes flow in a tube with a rough surface (see Figure 2). We found that for low to moderate roughness (the normalized standard deviation of the surface roughness below 5%), its effects on dispersion of a passive scalar are negligible, so that one can employ standard deterministic models that are much less computationally intensive.

References

[1] Xiu, D. and D. M. Tartakovsky, "Numerical methods for differential equations in random domains," *SIAM J. Sci. Comput.*, (under review) 2005.

[2] Tartakovsky, D. M. and D. Xiu, "Stochastic analysis of transport in tubes with rough walls," *J. Comp. Phys.*, (under review) 2005.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301) 903-5800
Gary.johnson@science.doe.gov

Automatic Differentiation of C and C++ Applications

B. Norris* (ANL), B. Winnicka (U. of Chicago and ANL), P. Hovland (ANL)

Summary

Derivatives, or sensitivities, are ubiquitous in scientific computing. They are used to solve inverse problems, including parameter identification and data assimilation. Derivatives are also used in numerical optimization and solution of nonlinear partial differential equations. Moreover, derivatives give an indication of the sensitivity of simulation outputs to changes in parameters, thus affording insight into physical processes. By providing accurate derivatives with minimal programming effort, automatic differentiation increases the productivity of computational scientists.

Automatic differentiation (AD) tools mechanize the process of developing code for the computation of derivatives. AD avoids the inaccuracies inherent in numerical approximations. Furthermore, sophisticated AD algorithms can often produce code that is more reliable and more efficient than code written by an expert programmer. ADIC is the first and only AD tool for C and C++ based on compiler technology. This compiler foundation makes possible analyses and optimizations not available in tools based on operator overloading. The earliest implementations of ADIC included support for ANSI C applications. Much more complete C++ coverage is available in our recent complete reimplementation (ADIC 2.0), which relies on EDG, a commercial C/C++ parser.

Component AD Tool Infrastructure

Modern AD tools, including ADIC, are implemented in a modular way (see Fig. 1), aiming to isolate language-dependent implementation features from the language-independent program analyses and semantic transformations. The component design leads to much higher implementation quality

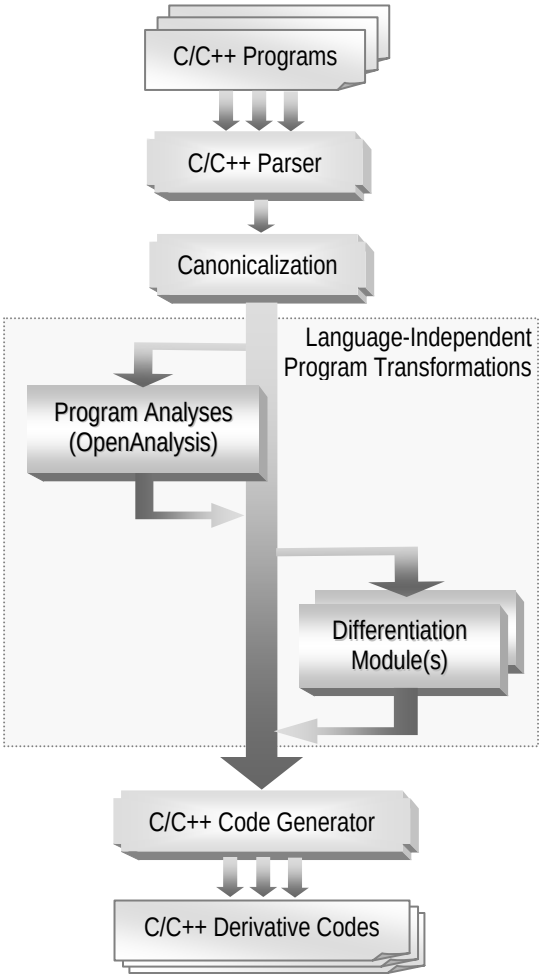

Figure 1. Overview of the AD process and component software infrastructure.

*(630)252-7908, Norris@mcs.anl.gov

because the different components can be implemented by experts in each of the different domains involved. For example, a compiler expert can focus on parsing, canonicalizing, and unparsing C and C++, while an expert in graph theory and algorithms can produce new differentiation modules without having to worry about the complexity of parsing and generating C++ code. This separation of concerns was achieved through the use of language-independent program analysis interfaces (in collaboration with researchers at Rice University) and a language-independent XML representation of the computational portions of programs. In addition to improved robustness and faster development times, this design naturally enables the reuse of program analysis algorithms and differentiation modules in compiler-based AD tools for other languages. In fact, the analysis and differentiation components are used in both ADIC and the OpenAD Fortran front-end (based on Rice's Open64 compiler).

Usability

As the name implies, AD is mostly automatic; however, some manual programming is involved in incorporating the automatically generated differentiated codes into an application. For example, the application developer must specify dependent and independent variables, initialize them correctly, and extract the results from the ADIC-generated derivative objects for further use in the computations. When one is working with arbitrary codes, these manual steps are unavoidable. When the application uses a numerical library, however, there arise opportunities for automating some of these steps. In collaboration with other researchers at Argonne, we have made the use of ADIC almost fully automatic in parallel applications that use certain PETSc features for the solution of nonlinear partial differential equations. Our long-term goal is to make the use of automatic differentiation for computing first and second-order derivatives virtually invisible to the user in as many numerical libraries as possible. We have made progress in that direction with the TAO (optimization) and PVODE (ordinary differential equations) toolkits.

ADIC Web Server

To make ADIC more accessible, we have provided a Web-based application server that allows users to invoke ADIC using any browser (see Fig. 2). The user can upload source code, select among various differentiation options, and invoke ADIC by a single button click. The differentiated sources can then be downloaded and compiled locally and linked against the small, portable runtime libraries.

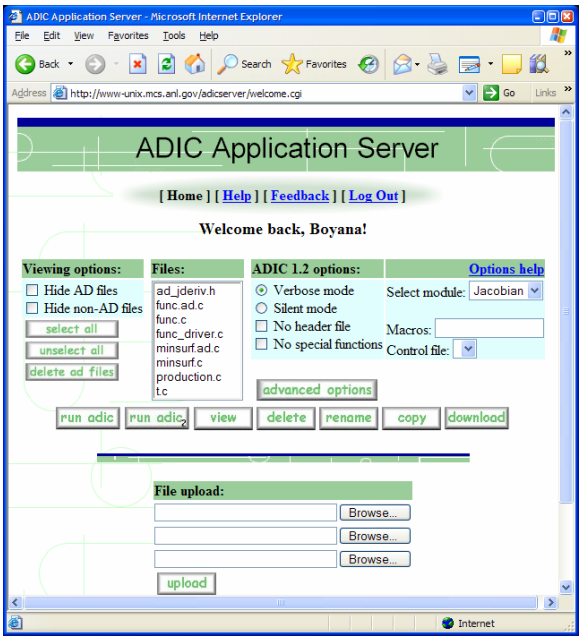

Figure 2. Screenshot of the ADIC Web-based application server.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division, Office of Advanced Scientific Computing Research, Phone: (301)903-5800
gary.johnson@science.doe.gov

Argonne Releases New Suite of Automated Deduction Tools

Larry Wos and William McCune* (ANL)

Summary

The automated deduction project at Argonne National Laboratory involves basic research in theory and methods, construction of production-quality software, and applications to problems in mathematics and logic. The goals of the project are to provide software and methods to assist engineers, scientists, and mathematicians with some of the deductive aspects of their work and to provide embedded reasoning engines for higher-level software. Accomplishments of the previous year include new methods that use known solutions to guide the search for solutions to related problems, new methods for refining existing proofs, a major software release, new applications to several areas of mathematics, and solutions to open problems.

Automated deduction problems are, in general, *undecidable.* That is, there exist no methods—even in theory—that are guaranteed to solve the problems. Therefore we have to *search* for solutions. But those search spaces are infinite, and hence many of the basic research questions focus on eliminating redundancy, restricting and directing the search, and developing effective inference rules.

The Power of the Abstract

As in most areas of mathematics and computer science, different classes of problem yield to different methods. As a result, substantial effort is given to developing methods for specific areas. We are, however, also devoting substantial effort to studying difficult problems in abstract algebra and logic, which expose core issues that are common to many other application areas.

Such abstract problems are usually easy to state but frequently very difficult to solve. Figure 1 contains three example theorems, first proved with Argonne's automated deduction systems (and not yet proved independently by human). The proof of XCB involves intermediate steps with thousands of symbols, and the proofs of MOL and Robbins have hundreds of complex equations.

Analysis of difficult proofs such as these has rarely been attempted, but we believe it can lead to valuable insights, for example, on the relationships between computer-generated and human-generated proofs. Therefore we have recently started a project to investigate the structure of complex proofs with the goal of developing more effective automated deduction methods.

*Mathematics and Computer Science Division, (630) 252-3065, mccune@mcs.anl.gov

Theorem XCB

Premises:
P(e(x,y)) & P(x) -> P(y)
P(e(x,e(e(e(x,y),e(z,y)),z)))
Conclusions:
P(e(e(x,y),e(y,x)))
P(e(e(x,y),e(e(y,z),e(x,z))))

Theorem MOL

Premise:
(yx)(((xx)z)(((((xy)z)z)x)(xu))) = x
Conclusions:
x((yz)(yz)) = y((xz)(xz))
x(y(x(zz))) = x(z(x(yy)))

Theorem Robbins

Premises:
x + y = y + x
(x + y) + z = x + (y + z)
((x + y)' + (x + y')')' = x
Conclusion:
(x' + y)' + (x' + y')' = x

Figure 1. Three example theorems proved with Argonne's automated deduction systems

A New Approach

A promising approach is to use semantics in the form of finite algebras to classify formulas in the searches and in the proofs. Using problems in Boolean algebra, lattice theory, ortholattices and subvarieties, cancellative semigroups, and nonassociative systems to study proof structure, we have obtained many new results in conjunction with the study of proof structure.

A New System

For the past decade, our primary production automated deduction system has been Otter. It has been used in a wide variety of applications, including basic research in abstract algebra and logic, systems modeling and verification, protocol analysis, natural language processing, and education. Otter has been used both as an embedded system and as a stand-alone system, both within Argonne and externally.

But the system has needed enhancements, and rather than continuing to make frequent patches and additions, we decided to develop an entirely new system, called ***Prover 9***.

Prover9, which will likely replace Otter for many applications, is an improvement over Otter in many ways.

- ***At the lower levels***, memory use is better organized giving better cache performance, memory consumption is lower, term indexing is faster, and basic deduction speed is greater.
- ***At the higher levels***, it has available more inference rules, more useful options for restricting and directing the search, and a much more effective *automatic* mode, by which the inference rules and search strategy are determined automatically by analysis of the problem or conjecture.

A New Suite of Tools

July 2005 marks the first release of ***LADR*** (Library for Automated Deduction Research), a new suite of automated deduction tools, including Prover9, which searches for proofs, and Mace4, which searches for counterexamples.

The software engineering in LADR is a great improvement over Otter's, enhancing maintenance, extensibility, and the modular construction of special-purpose deduction systems.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

NEOS: Network Enabled Optimization Services

Jason Sarich, Jorge Moré, and Todd Munson*
Argonne National Laboratory

Summary

The NEOS Project is a collaborative effort between Argonne National Laboratory and Northwestern University. The NEOS Guide provides educational material and case studies, while the NEOS Server gives users access to high-quality numerical optimization software through the Internet. The Server has been redesigned and reimplemented in the Python language to be more reliable, scalable, portable, and maintainable than earlier releases. An XML-RPC interface to the Server has been added to improve interoperability with other services.

The NEOS Server has been in existence for over ten years and now handles approximately 150,000 submissions per year from the academic, commercial, and government sectors. Figure 1 shows the growth in the number of submissions received by the Server. Over the years, the Server code had been extended by a variety of collaborators to provide additional features such as support for new clients, enhanced reliability, and improved scheduling. However, the code base had grown to such an extent that it was no longer practical to maintain.

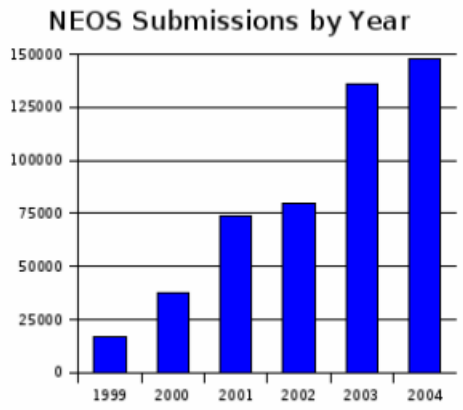

Figure 1. Number of submissions made to the NEOS Serve over the past four years.

Therefore, we redesigned and reimplemented the Server to improve maintainability and to reuse existing technologies without impacting reliability, scalability, and portability. The new version, NEOS Server 5, is now being extensively tested prior to replacing the current version. We highlight here some of the advantages of the new server.

Easier Maintenance. To implement the Server and communication package, we chose the Python programming language rather than Perl and tcl/tk. Python has many of the same advantages of Perl, for example, portability and extensibility, but also has an integrated object-oriented architecture, offers clean error handling, and is much easier to read and maintain. In particular, instead of requiring over 12,000 lines of Perl and 5,000 lines of tcl/tk, NEOS Version 5 requires only 7,500 lines of Python.

Improved Usability. The mechanisms used to represent, transfer, and store submissions to NEOS Version 5 were chosen to take advantage of existing standards and

* Mathematics and Computer Science Division, (630) 252-4279, tmunson@mcs.anl.gov

software. Instead of using homegrown parsing utilities, communication protocols, and database management, the new version now relies on XML to represent the data, XML-RPC for communications, and MySQL for database management. The reuse of existing software that is maintained and well documented not only allows us to concentrate our efforts on other needs but also provides a significant improvement in functionality, reliability, and usability over homegrown solutions.

The choice of XML-RPC, for example, gives users much more flexibility in accessing the NEOS Server from their own environments. Previously, the only published ways to access the NEOS Server were through Web Forms, E-mail, or the NEOS Submission Tool or from the AMPL or GAMS modeling languages by using the Kestrel clients we developed. Many users, however, have asked about accessing the NEOS Server from other environments, such as from a command line or inside a piece of software. Although technically possible, such an approach was cumbersome because it required users to write their own TCP/IP client socket code and crawl through the monolithic Perl and tcl/tk source code to find the required send and receive patterns. NEOS Version 5, on the other hand, has a published API (Application Programming Interface) callable through the XML-RPC protocol. In particular, users can now easily write their own clients that invoke the Server in a variety of programming languages by using XML-RPC.

Simplified Usage Monitoring. The improvements from using the MySQL database are generally seen not by the users sending problems to the Server but rather by the administrators. MySQL lets administrators easily monitor usage, quickly compile statistics, and flexibly mine the submissions based on different criteria. Further, it enables reliable tracking of available solvers, running processes, and jobs waiting to be scheduled. These features reduce the burden on the administrators to ensure that everything is configured correctly and running smoothly.

Better Use of Resources. The abundance of job submissions received by the NEOS Server means that we must make effective use of the available computational resources. In particular, a priority queue has been introduced in the scheduler to prevent the previously common occurrence of several long-running jobs exhausting the available resources to a point where all incoming jobs are backed up for hours or days. In particular, some of the computational resources are reserved for short-running jobs; any job designated as “short” is allocated five minutes of computational time on these resources before it is terminated. Moreover, by taking advantage of the usability and portability of the Python language, machines using the Windows operating system can now be used to process the submissions, potentially increasing the pool of available resources.

Our next step is to complete the beta testing of NEOS Version 5 and use it to replace the current Server. We then plan to concentrate on adding new features to further reduce the administration burden and implement new abilities based on feedback from our user base.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

SPARSKIT CCA Component

Masha Sosonkina*, Ames Laboratory

Summary

A set of CCA components has been developed for a package of iterative solutions of sparse linear systems. It is callable from a variety of programming languages and may be easily extended to incorporate novel solution techniques.

Under the Interoperable SPARSKIT Software project an extensible library of iterative solution techniques for general sparse matrices was created. It has been designed as a component of Common Component Architecture. The Babel interoperability language has been used to provide interfaces to SPARSKIT functionalities, such as preconditioners and accelerators. Due to Babel, the legacy Fortran 77 SPARSKIT code has been successfully integrated with newer additions, written in C, that have complex sparse matrix data structures. A coherent set of CCA components has emerged that is usable by client applications written in C, C++, and Python.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
gary.johnson@science.doe.gov

* (515) 294-6751, masha@scl.ameslab.gov

Multilevel Parallelism using Common Component Architecture and Global Arrays

Manojkumar Krishnan, Yuri Alexeev, Jarek Nieplocha, and Theresa L. Windus *,
Pacific Northwest National Laboratory

Summary

Multi-Component-Multi-Data (MCMD) is a novel direction in the Common Component Architecture (CCA) project. MCMD was shown to be effective for improving scalability of real scientific applications on large processor counts. When combined with Global Arrays (GA), a computational scheme that supports three different levels of parallelism was implemented in context of NWChem computational chemistry package. This approach, thanks to its ability to express variable parallel computational granularity, is expected to be critical for running complex scientific applications on future massively parallel systems.

The development of complex scientific applications for high-end systems is a challenging task. Addressing the complexity of software and algorithms involved is becoming increasingly difficult and requires appropriate software engineering approaches to address interoperability, maintenance, and software composition challenges. At the same time, scaling to massive numbers of processors is in itself a critical computational challenge. The requirements for performance and scalability to tens/hundreds of thousand processor configurations magnifies the level of difficulties facing the scientific programmer due to the variable levels of parallelism available in different algorithms or functional modules of the application.

Our recent paper accepted for SC'05, "Multilevel Parallelism in Computational Chemistry using Common Component Architecture," demonstrates how the CCA and GA can be used in context of computational chemistry to express and manage multilevel parallelism through the use of processor groups. Our target application is NWChem, a large (2.5 million lines of code) suite of computational chemistry algorithms that was developed based on multiple languages (Fortran, C, C++, Python) and programming models (MPI, Global Arrays). The experimental results for numerical Hessian calculation show that multilevel parallelism expressed and managed through the CCA component model can be very effective for improving performance and scalability of applications.

Multilevel Parallelism

We deployed multilevel parallelism by combining single-program-multi-data (SPMD) and multi-program-multi-data (MPMD) paradigms. The CCA framework by default supports SPMD/MPMD models. In the CCA framework, by default applications run as (SPMD), multi-component applications, meaning that each component is loaded across all of the processors. However, to create the dynamic environment to improve application efficiency and manage the resources effectively, multiple components (i.e. MCMD-component analog of MPMD) are instantiated on subgroups of processors using CCA *BuilderServices*.

* 509-376-4529, theresa.windus@pnl.gov

The development of multilevel parallel algorithms has been enabled by introduction of the processor group support in GA. One of the fundamental group-aware GA operations involves the ability to create shared arrays on subsets of processors. Every global array has only one associated processor group specifying the group that created the array.

In the GA programming model, data distributed in a processor group (containing *M* processors) can be redistributed to another processor group (containing *N* processors) regardless of the number of processors in each group and the data layout. This can be done as a collective call across processors in both the groups or as a non-collective one-sided operation. This feature enabled development of applications with nontrivial relationships between processor groups.

Example Application using Multilevel Parallelism: Numerical Hessian

In the overall algorithm, a CCA Hessian driver component is used to create subgroups to instantiate QM components to calculate gradients. These components in turn, may need to calculate numerical gradients and will therefore, need to create subgroups to calculate multiple energies. This particular algorithm offers essentially three levels of parallelism (Figure 1): one at the CCA (Hessian) level, one at the gradient level, and a level to calculate the energy (each energy itself can use a large number of processors).

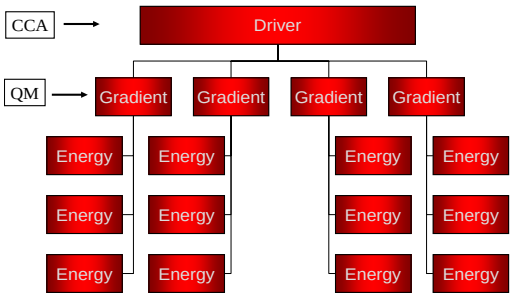

Figure 1. Example of subgroup use for numerical second derivative computations.

The scalability of numerical Hessian calculation is shown in Figure 2. The novel MCMD approach reduced the wall clock time by a factor of 10 on 256 CPUs.

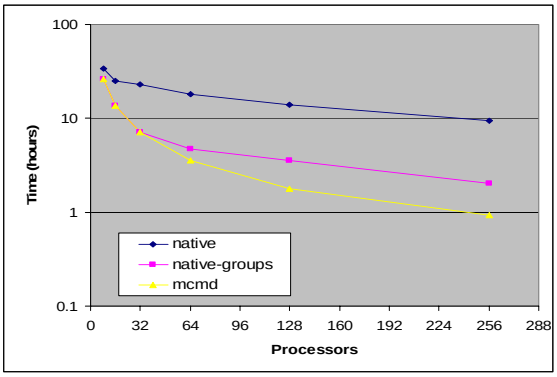

Figure 2. Scalability of Numerical Hessian on IA64 Linux cluster with Quadrics

Reusable Components in CCA Toolkit

A growing set of reusable high-performance computing components were developed at Pacific Northwest National Laboratory and included in the CCA toolkit this year. A distributed array component provides a uniform means to create and access distributed dense multi-dimensional arrays. A parallel I/O component provides simple interface to parallel I/O functionalities for scientific applications requiring out-of-core or check-pointing mechanisms. And a one-sided communication component provides remote memory access capabilities on various HPC interconnects such as Myrinet, Infiniband, and Quadrics. Also added were Lennard-Jones molecular dynamics, an application component for evaluating the performance of all the above components on massively parallel machines.

For further information on this subject contact:
Dr. Gary Johnson, Program Manager
Mathematical, Information, and Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: 301-903-5800
gary.johnson@science.doe.gov

TSTT "Terascale Simulation Tools and Technologies" Center

Harold Trease *, Pacific Northwest National Laboratory
J. Glimm, Brookhaven National Laboratory and
State University of New York at Stony Brook
D. Brown and L. Freitag, Lawrence Livermore National Laboratory

Summary

The Department of Energy (DOE) Scientific Discovery through Advanced Computing (SciDAC) Terascale Simulation Tools and Technologies (TSTT) Center is providing interoperable and interchangeable tools for advanced meshing and discretization application.

In partnership with Argonne National Laboratory (ANL), Brookhaven National Laboratory (BNL), Lawrence Livermore National Laboratory (LLNL), Oak Ridge National Laboratory (ORNL), Sandia National Laboratory (SNL), Pacific Northwest National Laboratory (PNNL), State University of New York (SUNY), Stony Brook, New York, and Rensselaer Polytechnic Institute (RPI), Troy, New York, the TSTT SciDAC project is chartered with making mesh-generation and discretization tools interoperable and interchangeable. The resultant advanced mesh-generation capabilities will be used within DOE SciDAC programs related to accelerators, fusion, climate modeling and biology.

The approach includes making DOE mesh-generation tools interoperable and interchangeable; using the DOE SciDAC CCA and SIDL/BABEL tools to achieve interoperability; and integrating TSTT mesh-generation and discretization technology into DOE SciDAC applications, such as accelerator design, fusion simulation, biology, and climate modeling.

Project accomplishments to date include a SIDL/BABEL mesh-generation interface that has been defined and is working. A discretization (mesh field-based) interface has also been defined, and it is being defined and worked on. TSTT mesh-generation tools are actively being integrated in to DOE SciDAC application areas. PNNL's application area is microbial cell biology, supporting the bioremediation of heavy metal waste products.

PNNL's focus is on meshing services for biology applications, and other TSTT collaborators are supplying similar technology to the areas of accelerator design for Stanford Linear Accelerator Center (LLNL, SNL and RPI), fusion (ORNL and SUNY), and climate modeling (ORNL and ANL).

PNNL scientists are applying TSTT tools to DOE bioremediation problems using the Virtual Microbial Cell Simulator (VMCS). The VMCS is one example where TSTT meshing and discretization technologies are being used successfully to construct a computational biology application. The main concept is to leverage this technology

* 509-375-2602, het@pnl.gov

to provide a general biological application tool by providing common interfaces and interoperability among a set of computational biology tools.

The main DOE/PNNL biology application area is microbial cell simulation applied to the bioremediation of heavy metal waste. The VMCS combines genome-scale metabolic network concept with a spatial, community model concept. Each individual microbe is modeled as its own genome scale metabolic network. These networks are coupled together in a self-organizing, spatial network, with communication between the organisms provided by multi-dimensional flow and transport models.

The following three figures describe the two major concepts behind this integrated biology application.

Figure 1. Different views of the Shewanella sp. *bacterium using scanning electron microscopy and electron tomography to obtain spatial information for constructing computational geometry and mesh models of the bacterium.*

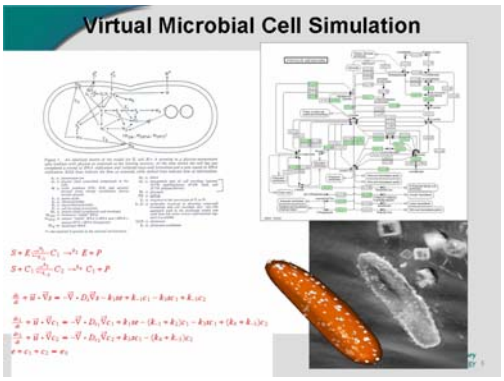

Figure 2. This figure demonstrates the coupling of four views of a microbial cell: The biologist's view (upper left), the genome view (upper right), the mathematical view (lower left) and the spatial view (lower right).

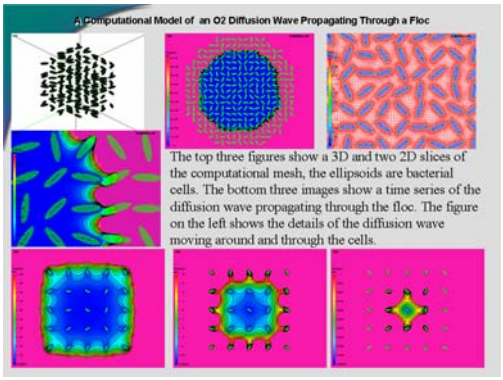

Figure 3. Demonstration simulation of the diffusion of oxygen through a "floc" of Shewanella sp. *bacteria.*

For further information on this subject contact:
Dr. Gary Johnson, Program Manager
Mathematical, Information, and Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: 301-903-5800
gary.johnson@science.doe.gov

Research on Scientific Components at Argonne

Lois Curfman McInnes* (PI), Steve Benson, Jay Larson,
Boyana Norris, and Jason Sarich
Mathematics and Computer Science Division
Argonne National Laboratory

Summary

The Center for Component Technology for Terascale Simulation Software (CCTTSS), a multi-institutional SciDAC-funded project, focuses on developing component technology for high-performance computing via the Common Component Architecture (CCA), with a goal of facilitating the interoperability and reuse of scientific software. Highlights of recent work by the Argonne contingent of the CCTTSS include developing scientific components, incorporating the CCA into climate simulations, improving the ease of use of CCA environments, and investigating issues in computational quality of service related to robust, efficient, and scalable performance.

Scientific Components. A key area of recent work among CCTTSS researchers at Argonne has been the development of production components that are used in scientific applications as well as prototype components that aid in teaching CCA concepts. Among the freely available components are tools for parallel "MxN" data redistribution, based on experience within the Model Coupling Toolkit, as well as tools for optimization and linear algebra. An important facet of this work is defining domain-specific common interfaces, which helps in realizing our vision of interchangeable scientific components. We are collaborating with the TOPS[1] SciDAC center to define common interfaces for linear and nonlinear solvers.

Optimization in Quantum Chemistry. An important goal of our project is to assure that scientific components are interoperable, able to deliver high performance, and useful to real applications. Thus, we are collaborating with chemists to perform CCA-based quantum chemistry simulations, which employ components based on the NWChem (PNNL) and MPQC (SNL) quantum chemistry codes for energy, gradient, and Hessian computations; the TaoSolver optimization component (ANL); and components based on Global Arrays (PNNL) and PETSc (ANL) for parallel linear algebra operations. Recent molecular geometry experiments have demonstrated reductions in simulation times up to 43% compared to the stand-alone chemistry packages.

Climate Components. Computational climate modeling is critical for our understanding of global processes and the consequences of human impact. Recent work at Argonne on component-based climate applications has leveraged the CCA's SIDL language-interoperability technology to allow access to the Fortran-based Model Coupling Toolkit (MCT) from C++ and Python. We have tested the C++

[1] *Terascale Optimal PDE Simulations Center*, PI: D. Keyes, http://tops-scidac.org

*Mathematics and Computer Science Division, (630) 252-5170, mcinnes@mcs.anl.gov

bindings in a parallel coupled application that uses the MCT's MxN and parallel interpolation facilities, while collaboration with an NSF-sponsored group at the University of Chicago has resulted in a Python parallel coupled test application. This work is one facet of building a proof-of-concept component version of the Community Climate System Model (CCSM). In addition, we are developing components for post-processing and analysis of climate simulations, including a component interface to netCDF, a popular data format in the geosciences.

Usability. Argonne has taken a leading role in a new effort focusing on the usability of high-performance component technologies. We have developed a new, fully automated build system for components and are implementing an Eclipse-based development environment that aims to provide automation whenever possible, while allowing the application developer to maintain control of the software development process.

Computational Quality of Service. Component-based environments provide opportunities to improve the performance and numerical accuracy of scientific simulations. The concept of the automatic selection and configuration of components to suit a particular computational purpose is called *computational quality of service (CQoS).* As part of our CQoS research, we have developed prototype software to support adaptive algorithms in component-based applications via runtime performance monitoring and performance database manipulations (see Figure 1). We have implemented components for manipulating two types of performance databases: a low-overhead runtime database for storing and accessing information about the currently executing application instance, and another, conventional database, which is accessed

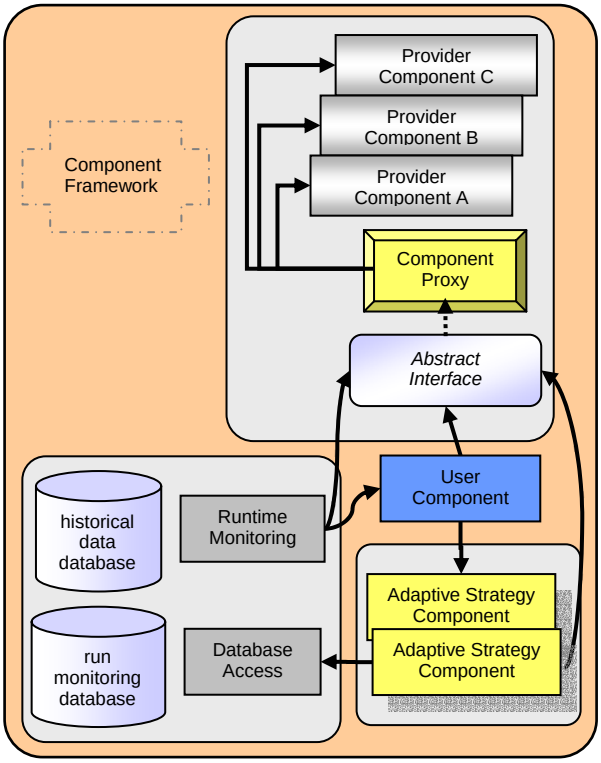

Figure 1. Computational quality of service infrastructure for adaptive solvers.

only at the beginning and end of an application to retrieve prior performance data and record any new information. The runtime database can be used for adaptive strategies, such as changing the linear solvers employed in the context of solving nonlinear partial differential equations. The larger and slower database containing information from multiple application instances is intended for offline analysis, in which large data sets can be processed and essential performance characteristics summarized.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

Science and Engineering Supported by PETSc

Satish Balay, William Gropp, Dinesh Kaushik, Mathew Knepley, Lois Curfman McInnes, Barry Smith* and Hong Zhang, Argonne National Laboratory

Summary

Partial differential equations (PDEs) are used to mathematically model phenomena in virtually all areas of science and engineering, from brain surgery to rocket science. We are developing the Portable, Extensible Toolkit for Scientific Computing (PETSc) to support high-performance simulations based on PDEs. The parallel computing infrastructure and scalable numerical solvers in PETSc enable scientists and engineers to focus on their primary scientific interests, thereby reducing implementation costs. PETSc helps hundreds of science and engineering groups to accomplish their work both faster and better.

Numerous ongoing application projects use various facets of PETSc. We highlight seven recent projects. All have resulted in scientific and engineering advances as chronicled in software and publications. References for these, as well as many others, may be found at our web site: www.mcs.anl.gov/petsc/petsc-as/publications

Fusion. Fusion uses the combination of hydrogen atoms to produce vast amounts of energy (for example, in our Sun) without the need for dangerous radioactive materials. Mathematical models for fusion are extremely complicated and difficult to solve but likely hold the key for learning to harness fusion energy for practical use.

Gyrokinetic simulation is one particular facet of computational fusion in which PETSc has been employed. This work involves the careful simulation of the dynamics of the electrons as they circle a toroidal chamber (the location where the fusion process is occurring). PETSc has also been used in other aspects of fusion simulation, including plasma models using both spectral and finite element methods.

Environmental Science. The Department of Energy (DOE) has an enormous interest in better understanding *subsurface flow*, the flow of water and other substances through the ground, for example, toward water wells. The reason is DOE is responsible for cleaning its legacy of contaminated sites across the country, a multibillion dollar endeavor. An example of this work is the LANL report *Modeling Thermal-Hydrological-Chemical (THC) Coupled Processes with Applications to Underground Nuclear Tests at the Nevada Test Site*, which is based on computations done using the PETSc software library.

Biological/Medical Applications. Advances in biological understanding and medical advances are proliferating as the direct result of mathematical models and numerical simulations. During the past year the PETSc software package and our expertise have been used for modeling and understanding *heart arrhythmias*. These disorders of the regular rhythmic beating of the heart cause the majority of sudden cardiac deaths.

Mathematics and Computer Science Division, (630) 252-9174, bsmith@mcs.anl.gov

Treating *fractures of the hip and spine* is a major medical cost. Understanding bone mechanical properties and how they may lead to fractures is thus an important medical research activity. A team of physicians and computational scientists have used PETSc and the software package Prometheus for simulations of the dynamics of hips and spines. This computation was a winner of a Gordon Bell Special Prize at SC2004 and ran scalably on over 4,000 processors.

MRIs and CT scans have had an enormous impact on the quality of medical care that is delivered today, but they still have limitations on the resolution of the images they can produce. Researchers are using PETSc to determine how to produce higher-quality imaging algorithms.

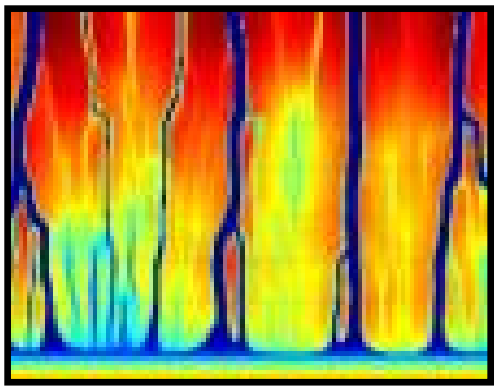

Fig. 1. Melt channels in magmatic reactive flow in porous ascent up a solubility gradient

Geosciences. The modeling of long-term *geodynamics* is crucial, not only for energy exploration and exploitation, but for a more fundamental understanding of processes involved. Researchers at Columbia University have used PETSc to explore the coupling between regional mantle subduction and the reactive flow of hydrous magma in the lithosphere; see Fig. 1.

A key component to the reactive flow simulation is the introduction of a semi-Lagrangian advection solver. By eliminating numerical diffusion, we were able to resolve the advective dynamics; see Fig. 2. This functionality has been packaged as a component and incorporated into PETSc, allowing any field to be advected precisely.

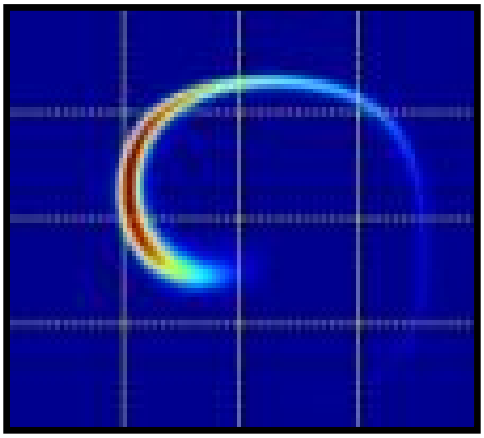

Fig. 2. Sharp features can easily be resolved using semi-Lagrangian advection solvers

Ship-Sea Interaction Modeling. The U.S. Navy has a strong interest in the simulation of ships and submarines moving through the ocean. The software package CFDShip-Iowa was recently enhanced to use PETSc for its algebraic solver needs. This package is a multiblock free surface flow solver that uses overset structured meshes, designed for ship applications. It has been used in the past year for "roll for surface combatant," for "beveled-trailing edge flows and wakes," and for "non-body-of-revolution" submarines as well as other applications.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
gary.johnson@science.doe.gov

Advanced Computational Modeling for Reactor Design

Matthew G. Knepley* and Dmitry A. Karpeev, Argonne National Laboratory

Summary

A variety of efforts at Argonne support and enable scientific research through mathematical modeling and simulation. These efforts leverage the substantial analytical expertise and the advanced computational technology developed at Argonne and allow researchers to probe regimes inaccessible to traditional experimental techniques. Recently these efforts have included the development of scalable parallel simulations in support of the design of the Very High Temperature Reactor being developed as a vehicle for clean hydrogen production on the large scale.

The Portable, Extensible Toolkit for Scientific computing (PETSc), developed at Argonne, is the premier scalable parallel numerical library, used in dozens of scientific applications around the world. PETSc helps bring the power of DOE's most advanced computational capabilities to both new and established scientific simulation codes. Key to PETSc's design is its encapsulation of the complexity of the underlying parallel numerical algorithms and its presentation of them in terms of the mathematical abstractions most familiar to scientists.

Recently PETSc was successfully applied in support of the design of the Very High Temperature Reactor at the Idaho National Laboratory (INL). Specifically, PETSc was used to extend the fully compressible flow codes, most notably the pressure-corrected implicit continuous Eulerian scheme (PCICE) developed at INL. Figure 1 illustrates the capabilities of PCICE at highly accurate resolution of compressible flows and heat transfer in complicated geometries. While PCICE is a state-of-the-art code capable of very accurate resolution of the wide range of fluid flows encountered in the core of a pebble-bed reactor, it lacks the parallel capabilities that would enable large-scale simulations. We have used PETSc to parallelize PCICE with minimal impact on the original code. Moreover, PCICE now has access to a variety of scalable linear solvers and flexible mesh manipulation tools offered by PETSc.

Extending PETSc's capabilities. Some of the capabilities introduced in our work with

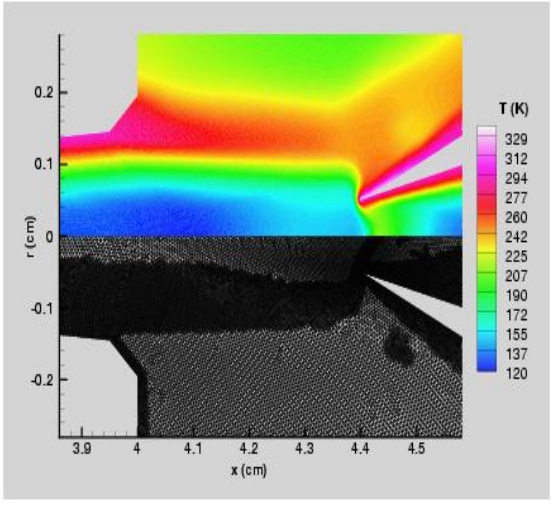

Figure 1. Temperature distribution obtained with PCICE in a nozzle subject to flow through it. Courtesy of Richard Martineau, INL.

*Mathematics and Computer Science Division, (630)252-1870, knepley@mcs.anl.gov

PCICE have necessitated extensions of PETSc to enable the handling of complicated geometric data, such as the structure of computational meshes and fields over them. To approach the additional complexity in a uniform way, we have introduced the notion of a *Sieve,* which abstracts the structure of topological objects, such as meshes, and extends it to more general hierarchical structures that can benefit from a geometric viewpoint.

The essential idea behind a Sieve is to consider a geometric object as composed, or covered by, more elementary subobjects. The utility of Sieves lies in their ability to manipulate distributed decompositions of arbitrary complexity in a scalable way presented in a concise interface. The numerical data attached to Sieves, such as the pressure and temperature fields defined over computational meshes, are regarded as the *sections* of a *fiber bundle* over the Sieve, to borrow a term from differential geometry. The analogy is more than just terminological: the decomposition of a Sieve allows us to define the restriction and prolongation operations on the data, so that computations can proceed in the local setting of a subobject, and later be automatically reassembled.

Applications. We anticipate that the newly introduced capabilities will help, not only in the scalable parallelization of PCICE, but also in its coupling to the EVENT neutronics code from Georgia Tech, completing an essential kernel of a reactor simulation. Further advantages include the ability to implement flexible parallel adaptive mesh refinement and load-balancing strategies using only a serial mesh generator and the mesh improvement algorithms also developed at Argonne.

In addition to reactor design applications we envision using Sieves and PETSc's geometric capabilities in general to analyze the structure of large genetic and metabolic networks. We anticipate the development of metabolic engineering tools for the maximization of flux through critical metabolic pathways, such as those responsible for the hydrogen production in bacteria.

Furthermore, scalable parallel codes for hierarchical nonlocal field computations, such as those encountered in micromagnetics and electrostatics, will greatly benefit from the powerful topological abstractions. They are based on multipole expansions or the interaction kernel interpolation, whose implementation can be simplified into many cooperating Sieves. Moreover, the nontrivial global topology of micromagnetic formulations is also easy to incorporate.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
gary.johnson@science.doe.gov

Scalable Solvers – *Toolkit for Advanced Optimization*

Steven Benson*, Jorge Moré, Lois McInnes, and Jason Sarich
Argonne National Laboratory

Summary

Discretizations of continuous optimization problems often lead to nonlinear optimization problems with many degrees of freedom. The Toolkit for Advanced Optimization (TAO) enables scientists to apply state-of-the-art optimization algorithms, preconditioned linear equation solvers, automatic differentiation, multilevel methods, and parallel hardware to solve these problems and achieve scalable performance.

Examples of continuous optimization problems include the variational form of elliptic partial differential equations and obstacle problems from physics and engineering. Obstacle problems, as shown in Figures 1 and 2, minimize the area of a surface fixed at the boundary and stretched over obstacles. Finite element methods approximate these problems by creating a mesh, discretizing the domain, and formulating the optimization problem with a finite number of variables. However, the number of variables required for a sufficient approximation may still be very large and demand computationally intensive methods.

At Argonne we are exploring scalable algorithms for mesh-based optimization problems in the sense that the number of operations required to solve the problem grows linearly with the number of variables: this is an ambitious requirement from a computational viewpoint. Scalable algorithms for mesh-based problems generally required mesh-invariance in the sense that the number of iterations is independent of the granularity of the mesh. In theory, mesh-invariance can be obtained with a traditional Newton's method, but in practice numerous complications arise. The main difficulties are the generation of the mesh, the computation of derivatives, and the preconditioning of the Hessian matrix.

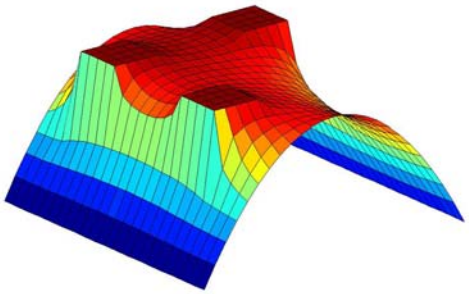

Figure 1: Obstacle problem on a fine mesh.

The lack of efficient numerical software makes the development of scalable algorithms in a single-processor environment difficult. Parallel-processor environments magnify these difficulties because the overhead of communication between processors must be balanced with other performance considerations. As a result, many practitioners are left to assemble their own methods. To address these problems, the Toolkit for Advanced Optimization (TAO) has implemented

* Mathematics and Computer Science Division, (630)-252-7232, benson@mcs.anl.gov

efficient optimization algorithms and coupled them with two other software packages developed at Argonne. TAO version 1.8, released in May 2005, uses ADIC to compute derivatives and PETSc to precondition the linear equations. By partitioning the domain over multiple processors, users of TAO can solve mesh-based problems in parallel and achieve scalable performance.

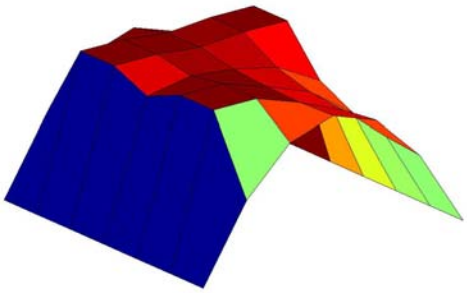

Figure 2: Obstacle problem on a coarse mesh can be coupled with a fine mesh to provide scalable performance in the optimization solvers.

Newton's algorithm for the optimization problem requires the first- and second-derivative information from the objective function. Writing code that evaluates the function can be difficult and prone to error. Automatic differentiation is a technique for generating code that computes gradients and higher-order derivatives. Given code that computes the objective function, this software generates additional code that computes the function and its gradient. Using TAO, scientists need only implement the objective function over a single element with a few variables; the toolkit will apply ADIC, the automatic differentiation tool developed at Argonne, to compute the derivatives in parallel.

The performance of optimization solvers depends heavily on the performance of the linear solver in the algorithm. TAO uses PETSc, also developed at Argonne, to apply Krylov methods with a variety of preconditioners such as incomplete LU factorization and additive Schwartz methods.

Mesh sequencing is a technique for solving mesh-based problems that uses the solution of a problem on coarse mesh solution as the initial starting point for a finer mesh. This multilevel method is a standard technique for solving systems of nonlinear differential equations, but with few exceptions, it has not been used to solve mesh-based optimization problems.

Table 1: Performance of TAO Solver on Obstacle Problem on 16 Processors.

Mesh Variables	Iterations	Seconds
1121 x 1121	4	2.5
2241 x 2241	1	2.9
4481 x 4481	1	9.4
8961 x 8961	1	29.4

Table 1 shows the time needed to solve the obstacle problem on four distinct meshes. The numbers reflect the use of a Newton solver with an incomplete LU preconditioner, ADIC, and mesh sequencing. The computations used 16 processors from the "Jazz" Linux cluster at Argonne's Laboratory Computing Resource Center. The data shows that as the number of variables increases, the time required to solve the problem increases proportionately. TAO solvers are developed with this kind of scalability in mind.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

Scalable Linear Solvers

Robert Falgout*, Van Henson, and Panayot Vassilevski
Lawrence Livermore National Laboratory

Summary

The goal of this project is to develop parallel multigrid methods for solving the large linear systems of equations that arise in many DOE scientific simulation codes. These methods have the potential to reduce simulation times dramatically (by as much as a factor of 10 or more), enabling new advancements in science. This project focuses on the longer-term fundamental mathematics of multigrid methods research and aims to make the method applicable to a broader class of applications.

At the core of many DOE scientific simulation codes is the need to solve huge linear systems on thousands of processors. Multigrid methods are so-called *scalable* or *optimal* methods because they can solve a linear system with N unknowns with only $O(N)$ work. This property makes it possible to solve ever larger problems on proportionally larger parallel machines in constant time. Classical iterative methods like conjugate gradients are not scalable and can be orders of magnitude slower than multigrid methods on the large problems of interest here.

Multigrid works because simple iterative methods (like Jacobi or Gauss-Seidel) are effective at reducing high-frequency error, while the remaining low-frequency error can be accurately represented and efficiently eliminated on coarser grids. Applying this simple multigrid idea to get a scalable method often involves considerable algorithmic research. One has to decide which iterative method to use as a smoother, how to coarsen the problem, and how to transfer information between the grids. When designed properly, a multigrid solver will uniformly damp all error frequencies with a computational cost that depends only linearly on the problem size. In all multigrid methods the smoother and coarse-grid correction steps must complement each other well.

The algebraic multigrid (AMG) method was originally developed to solve general matrix equations using multigrid principles. The fact that it used only information in the underlying matrix made it attractive as a potential black box solver. Instead, however, a wide variety of AMG algorithms have been developed that target different problem classes and have different robustness and efficiency properties.

The goal of our research is to develop new AMG algorithms and techniques that broaden its range of applicability even further. In AMG, the smoother is typically fixed, so most of the research focuses on how to build the coarse-grid correction.

* Center for Applied Scientific Computing, (925) 422-4377, rfalgout@llnl.gov

One model success story this past year was the application of our Adaptive Smoothed Aggregation (αSA) method to Quantum Chromodynamics (QCD, currently funded under SciDAC) applications. The αSA method is the only method ever to exhibit scalable convergence behavior on a QCD application independent of physics and discretization parameters, and αSA was shown to be faster than existing methods even on today's relatively small simulations.

Nearly all multigrid methods today require some a-priori information about the near null-space of the operator to be effective. However, for the QCD application, the near null-space of the operator is not known a-priori. Furthermore, these components can be oscillatory (see Figure). The αSA method automatically uncovers these problematic components and makes adjustments for them (i.e., it adapts).

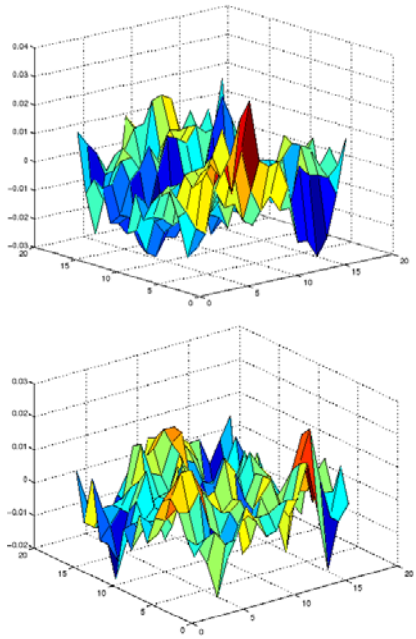

Figure 1: Real (top) and imaginary (bottom) parts of the near null-space in QCD are not known a-priori and are oscillatory. Our αSA method automatically uncovers these problematic modes and adjusts to them.

There were several new additional solver developments this past year that we hope will have a similar impact on DOE applications in the future. We have developed a new variant of the adaptive multigrid idea and shown it to be robust for difficult PDEs with large near null-spaces such as Maxwell's Equations on unstructured grids. Also, based on our new sharp convergence theory, we have developed a more predictive form of compatible relaxation (CR) that is aimed at further improving the robustness of AMG. CR is an efficient technique for measuring the quality of coarse grids, and hence, can be used to select coarse grids in AMG.

The co-PIs were invited to give eight talks at prestigious conferences, including a plenary talk on adaptive AMG at the European Multigrid Conference, and a Topical Lecture on AMG at this year's Society of Industrial and Applied Mathematics (SIAM) Annual Meeting. Our paper on αSA was selected to appear as a SIGEST article in SIAM Review. Two additional journal articles have been submitted for review and another one is to appear.

The αSA algorithm was developed in collaboration with researchers at the University of Colorado at Boulder, and the sharp convergence theory was developed with Penn State. See our publications at http://www.llnl.gov/CASC/linear_solvers/.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

New Multilevel Solvers for Computational Fluid Dynamics

R. Tuminaro*, M. Gee, J. Hu, M. Sala, J. Shadid
Sandia National Laboratories

Sandia researchers have developed two new linear solver algorithms, which will improve the performance of many fluid flow simulations central to the advanced modeling performed at the national laboratories. Among the applications that currently use solvers we have developed are chemical vapor deposition computations for semiconductor and MEMS processing, aerodynamics calculations to determining flight characteristics, and combustion simulations to understand pool fires and validate weapon systems (see [2-5] for recent fluid flow/solver studies). Some of these simulations have been performed on equation sets with over 100 million unknowns on thousands of processors. Advances in linear solver technology translate directly into increased capabilities within these applications as the linear solution is typically one of the primary computational bottlenecks. These increased capabilities lead to more detailed and physically realistic studies over larger ranges of physical and temporal scales.

In the area of incompressible fluid flow, we have developed a new scalable block preconditioner (S-LSC) that uses multigrid solvers for the sub-blocks. S-LSC exploits approximate commuting properties of differential operators [1] and represents a substantial improvement over existing methods within the family of pressure convection-diffusion solvers. The key idea in developing S-LSC involves the formulation of a least squares algebraic commutator equation that is combined with stabilization techniques for discretizations. We have routinely demonstrated several factors of improvement using S-LSC over the previous best methods within this class. More importantly, S-LSC requires significantly less user interaction than the previously most popular method (the *Fp* scheme). This makes S-LSC significantly more suitable for use by non-experts in linear solvers and much more practical to integrate into complex applications codes found within the laboratories. The first version of our algorithm (LSC) is described in a recent book (Elman, Silvester Wathen, *Finite Elements and Fast Iterative Solvers***,** Oxford Univ. Press'05). It is also incorporated into a research code distributed by the University of Maryland, IFISS, (http://www.cs.umd.edu/~elman/ifiss.html) and will be implemented in the Trilinos package (http://software.sandia.gov/trilinos). Our more recent extension of LSC to stabilized discretizations (S-LSC) will likely appear in future versions of the codes and book mentioned above. Table 1 illustrates typical results obtained with S-LSC.

Re	**Fp**	**S-LSC**
10	37	22
100	79	25
200	100	27

Table 1: Iterations vs. Reynolds number on a backward facing step with Q1-P0 elements within a Newton solver.

* 925-294-2564, tuminaro@ca.sandia.gov

We have also recently devised a new algebraic multigrid method for compressible flow based on a generalization of the highly successful smoothed aggregation method for symmetric systems. The initial intended application of this new technique (called NSA) is within a compressible flow code, PREMO, to simulate the aerothermodynamics of weapon systems. The recently devised method required several new concepts to be developed including a suitable notion of *energy* and *local* energy minimization. Several problems have been addressed exhibiting significantly improved convergence rates over existing smoothed aggregation algorithms. In several cases the more standard smoothed aggregation methods diverge while the new scheme attains a solution without difficulty. The convergence history for a single linear solve within a sample calculation is depicted below and illustrates a reduction in iterations of more than a factor of two compared to plain aggregation on a Falcon jet calculation (shown below).

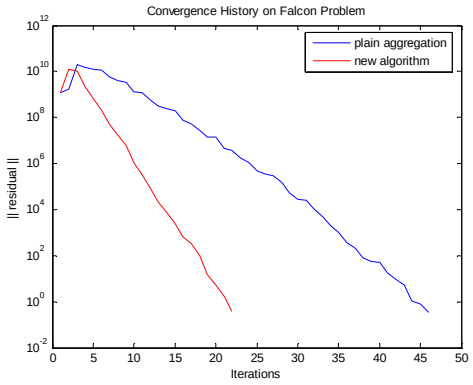

Table 2 illustrates typical results for a model convection-diffusion problem. To the best of our knowledge this is the first truly promising generalization of smoothed aggregation to non-symmetric systems.

ε	*NSA*	$\omega = \varpi = 0$	SA
10^{-2}	7	54	7
10^{-3}	8	60	NC
10^{-4}	16	59	NC

Table 2: Iterations to solve on a square for $\varepsilon\Delta u + 4x(x-1)(1-2y)u_x + 4y(y-1)(1-2x)u_y = rhs$.

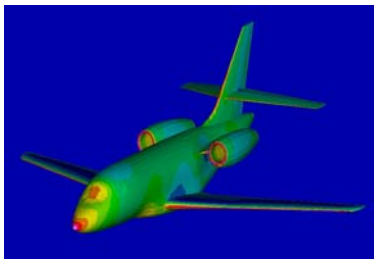

[1] *Elman, Howle, Shadid, Shuttleworth & Tuminaro.* ***Block Preconditioners Based on Approximate Commutators****, SISC'*05.

[2] *Lin, Sala, Shadid, & Tuminaro.****Performance of Fully-Coupled Algebraic Multilevel Domain Decomposition Preconditioners for Incompressible Flow and Transport, ion of the Generalized Global Basis Method for Nonlinear Problems****, submitted to IJNME'05.*

[3] *Shadid, Salinger, Pawlowski, Lin, Hennigan, Tuminaro & Lehoucq.* ***Large-Scale Stabilized FE Computational Analysis of Nonlinear Steady State Transport/Reaction Systems,*** *CMAME'05.*

[4] *Shadid, Tuminaro, Devine, Hennigan & Lin.* ***Performance of Fully-Coupled Domain Decomposition Preconditioners for Finite Element Transport/Reaction Simulations****, JCP, Vol 205, 1, 24-47,* 2005.

[5] *Waisman, Fish, Tuminaro & Shadid.* ***Acceleration of the Generalized Global Basis Method for Nonlinear Problems****, JCP'*05.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
gary.johnson@science.doe.gov

High-Order Methods for MHD

Paul Fischer*, Argonne National Laboratory

Summary

Magnetohydrodynamics (MHD) is central to many challenging physics problems of importance to the DOE Office of Science. We have developed a state-of-the-art code that features high-order numerical discretizations and multigrid solvers capable of scaling to thousands of processors to simulate MHD in complex domains.

MHD governs the motion and stability of many physical phenomena of interest to DOE, including astrophysical plasmas, geo- and solar dynamos, fusion plasmas, liquid metal cooling systems in nuclear reactors, and liquid metal plasma-facing material in tokamak side-walls and diverters. Most of these applications are dominated by convective transport and operate at high hydrodynamic and magnetic Reynolds numbers, Re and Rm, respectively. The governing physics is thus highly nonlinear and is essentially nondissipative. In addition, many of these applications involve complex domains that preclude the use of traditional global spectral methods.

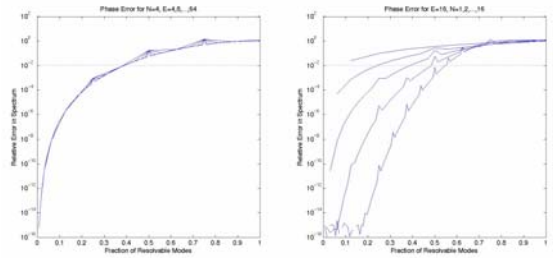

Figure 1. Phase error vs. fraction of resolvable modes for h-refinement (left) and p-refinement (right) for $u_t + u_x = 0$. The fraction of resolvable modes is increased only with increased order (p-refinement), which also yields rapid error reduction.

Working with University of Chicago astrophysicists Fausto Cattaneo and Aleks Obabko, Argonne researchers have developed a numerical code for simulation of liquid metal MHD that couples state-of-the-art high-order numerical methods with the geometric flexibility required for the challenging MHD problems facing the DOE and scientific community.

Our MHD code is based on Argonne's hydrodynamics code Nek5000, which is a past Gordon Bell Prize winner that readily scales to thousands of processors. Nek5000 is based on the spectral element method (SEM), a high-order weighted residual technique that combines the geometric flexibility of finite elements with the rapid convegence and tensor-product efficiencies of global spectral methods. As with the finite element method, functions in the SEM are represented on compactly supported subdomains (elements), thereby simplifying implementation of complex boundary conditions. Grid refinement in the SEM is achieved by increasing the order of the polynomial representation within each element, with typical orders in the range of N= 8–16. The use of such high order minimizes numerical dissipation and dispersion and is important for high-Reynolds number applications where high

* Mathematics and Computer Science Division, Argonne National Laboratory, (630) 252-6018, fischer@mcs.anl.gov

wave-number error components are only weakly damped by physical viscosity. This point is illustrated in Fig. 1, which shows that increasing approximation order leads to efficient use of computational resources. For modest error tolerances, a fivefold reduction in the number of gridpoints *per space dimension* is achievable by going from linear elements to 12^{th}-order elements.

In addition to minimizing MHD discretization errors, we have made significant strides in stabilized methods, on multigrid solvers, and porting to terascale platforms. Of particular importance to MHD is the development of dealiased quadrature rules that ensure energy conservation (Fig. 2, left). We have recently developed spectral element multigrid techniques that have proven to be two to three times faster than our earlier multilevel Schwarz methods across a range of applications. We have ported the new MHD code to NERSC's Seaborg platform (for which we received a two million node-hour allocation as part of a DOE 2005 INCITE Award) and to Argonne's 2048-processor IBM BG/L. Nek5000 is currently sustaining > .9 Tflops on BG/L. We expect this performance to improve significantly with recently developed matrix-matrix product kernels that make better utilization of the BG/L dual-core processors.

We are working with experimentalists led by Hantao Ji at Princeton Plasma Physics Laboratory (PPPL) to study the magneto-rotational instabilities that are believed to be responsible for the generation of turbulence in (magnetized) accretion disks. (The gravitational energy released by accretion—the in-fall of material into a central potential well—is believed to power many of the energetic phenomena observed in the universe.) Rotational flows that are nominally stable can become destabilized through the introduction of a weak magnetic field. Our initial simulations have demonstrated that this is indeed the case for a particular choice of Re and Rm. Because the ratio Rm/Re is a fixed property of the fluid medium, experiments are limited to relatively small values of Rm and thus can not reach the highly nonlinear magnetic regime. The computations are restricted to smaller values of Re, but are able to simulate at much larger values of Rm. The numerics and experiments are thus complementary and will be able to map out a larger region of the Re – Rm parameter space than would be possible using a single mode of investigation.

We have completed a sequence of axisymmetric hydrodynamic and MHD simulations and have initiated a sequence of hydrodynamic simulations (Fig. 2, right) that will make allow comparison with the PPPL data. The PPPL collaboration has led to a new DOE funded project on the effects of MHD on the stability of free-surface flows.

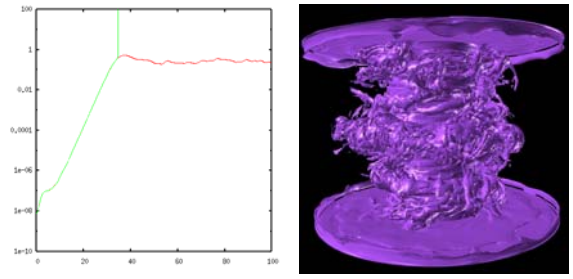

Figure 2: (left) History of magnetic energy in an MHD benchmark for aliased (green) and dealiased (red) nonlinear evaluations. In the saturated nonlinear state, the aliased case is numerically unstable. (right) Simulation of a turbulent flow-field for the PPPL experimental configuration.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

High-Order Methods for Nanophotonics

Misun Min* and Paul Fischer, Argonne National Laboratory

Summary

Argonne researchers have recently developed high-order methods in electromagnetics that demonstrate exponentially accurate computational results. Further development of the methods and successful extension to the application in nanophotonic waveguides and photonic crystals will significantly increase the importance of higher-order methods and their use by the optics community.

In the study of light interacting with a metallic nanoscale object, a particular computational issue is that the problem includes sharp discontinuities in the dielectric function along the surface of the metallic object. In such cases, standard lower-order methods require considerable computational work in order to achieve a certain expected accuracy. The drawback stems from the slow rate of convergence of the methods for discontinuous problems or problems whose solutions have less regularity in smoothness.

Argonne researchers have been developing and analyzing efficient and accurate high-order numerical methods for the study of electromagnetic wave propagations for nanophotonic and photonic crystal waveguide problems. Three high-order numerical techniques have been developed for those problems: pseudo-spectral Fourier time-domain method combined with postprocessing techniques, spectral-element methods (SEM), and discontinuous Galerkin spectral element methods (DG-SEM).

We applied the standard pseudo-spectral Fourier time-domain method (PSTD), combined with postprocessing techniques such as Gegenbauer and Pade´ reconstruction, to nanophotonic problems. Since pseudo-spectral methods are global methods, numerical artifacts such as high oscillations (Gibb's phenomena) deteriorate the approximate solutions because of the nonsmoothness between materials. Careful numerical analysis, however, reveals that the PSTD method retains accurate phase information despite the oscillatory nature of standard Fourier reconstructions. Gegenbauer postprocessing, applied at a final time step to PSTD solutions, successfully enhances their accuracy and provides an accurate reconstruction.

Figure 1 shows the original (top left) and Gegenbauer reconstructed (top right) results based on 512 x 512 PSTD data. The Gegenbauer reconstruction eliminates the nonphysical oscillations of the original PSTD calculation. The reconstructed results agree well with the finer resolution finite difference time domain (bottom) results away from the metal boundary. Moreover, near the material interface, Gegenbauer-postprocessed results successfully capture a more physically reasonable profile.

* Mathematics and Computer Science Division, (630) 252-5380, mmin@mcs.anl.gov

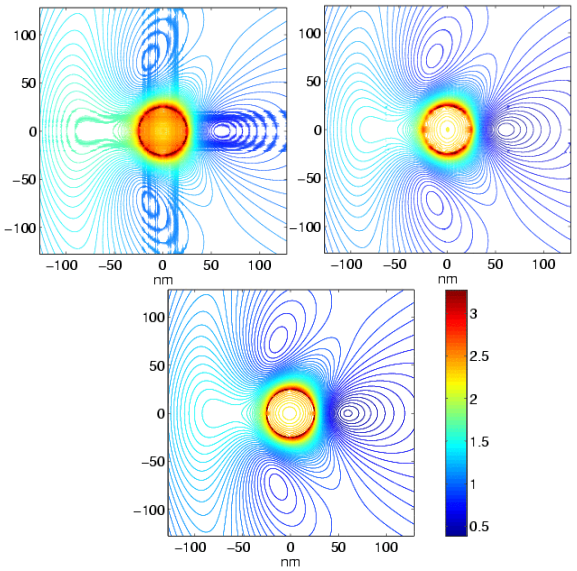

Figure 1. Incident plane-wave field response computed with PSTD (top left), PSTD + Gegenbauer reconstruction (top right), and finite differences (bottom). Resolution is 512 x 512 (top) and 2560 x 2560 (bottom).

We have also demonstrated that DG-SEM and SEM are very effective for simulating propagation of electromagnetic waves in free space and for computing Maxwell's eigenvalue problems in one and two dimensions. We split the domain into subdomains in a way that the dielectric function in each subdomain is smooth. This approach reformulates the problem as continuous in each subdomain, and we use high-order spectral bases within each subdomain. These methods retain many of the advantages of the PSTD method (principally, minimal numerical dispersion), while avoiding the Gibb's phenomenon entirely. Contrary to conventional low-order methods that give generally second-order or at most third-order convergence, the spectral-element method can dramatically reduce the computational cost in two- and three-dimensional problems with high accuracy.

We also have developed discontinuous Galerkin spectral element codes for electromagnetics problems incorporated into Argonne's spectral-element code Nek5000, which is recognized for its algorithm quality and scalability. With these codes, we are carrying out time-domain simulations of propagating plane waves impinging on a nanocylinder. A criticial development for the simulation of practical problems has been the incorporation of perfectly-matched-layer boundary conditions that admit reflectionless outgoing waves.

In the coming year, we plan to extend this work to two- and three-dimensional studies of nanophotonic waveguides for various structures and to periodically structured photonic crystals.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

Multiscale Density Functional Theory and Computations

Enabling system-level simulations of nanoscale systems

Mihai Anitescu*, Dan Negrut, and Todd Munson, Argonne National Laboratory

Summary

Density functional representation is the basic modeling tool of nanoscience. Solving the optimization problem associated with this representation is the key to predicting the chemical and mechanical properties of nanoscale systems. We use multiscale mathematics to expand the reach of computational approaches from the thousands-of-atoms configurations that they can currently solve to the millions-of-atoms configurations that they need to solve.

Nanoscience, the science related to the understanding and control of the matter at the nanoscale, the length scale of 1–100 nm, carries the promise of significantly improving our everyday life, in the immediate and medium-term future. Using nanosciene to obtain novel nanocatalysts, for example, will result in far-improved catalytic converters and substantially reduced pollution from our cars. The all-important aim of clean energy depends heavily on the ability of efficiently storing hydrogen storage, for which nanoscience is expected to provide a broad range of solutions.

Nanoscience allows access to a vastly enhanced array of new material properties. The electronic structure of materials undergoes dramatic changes when their dimensions are reduced to the nanoscale regime, leading to new regimes of physical, mechanical, and chemical behavior not observed in bulk materials. For example, nanoscale oxide structures exhibit unusual characteristics compared with their bulk counterparts: thermodynamic stability and lattice properties, magnetic properties, ion transport, optical properties, chemical reactivity, and ferroelectric properties. Nanoscale metals, polymers, and semiconductor materials have also been shown to exhibit unique properties and a wide range of technological applications .

A key endeavor in unlocking the potential of nanoscience is the ability to use high-performance computing to predict the behavior of nanoscale materials to a high accuracy. When the chemical properties of such materials are sought, as is the case in virtually all energy-related applications, the distribution of electrons inside of the material must also be correctly determined. Therefore, any computation must accurately solve the density functional theory (DFT) problem whose result is the distribution of electrons (the electron density) at the surface and inside the nanoscale material. For large nanoscale systems (that may have millions of atoms), such computations are outside the reach of stat-of-the-art computational techniques for DFT. Such techniques currently take several weeks on 64-

* Mathematics and Computer Science Division, (630) 252 4172, anitescu@mcs.anl.gov

processor machines for configurations with a few thousand atoms.

To remove this bottleneck, we developed a multiscale approach that computes the electron density in the nanoscale system, while carrying out detailed computations only in small subdomains, each of which is of a size that is approachable by current DFT methods. The aggregate number of atoms in the subdomains is only a fraction of the total number of atoms, which makes the overall computation within reach in the immediate future. In Figure 1, we present a comparison between the multiscale approach and the direct simulation approach for a one-dimensional, Thomas-Fermi DFT problem. We see that the multiscale approach (interpolation reconstruction) produces virtually indistinguishable results while using less than half the number of degrees of freedom (mesh points) to carry out the computations. The results validate our rigorous mathematical analysis of the approach.

We are currently investigating extensions of our technique to generic DFT approaches, such as the highly accurate Kohn-Sham approach. In addition, we are pursuing advanced techniques for efficiently solving the resulting very large scale optimization problem. To harness the potential of our advances, we are building the three-dimensional multiscale DFT code MS-DFT. Computations of the forces on the nuclei in a 1000-atom subdomain of the nanoscale material is presented in Figure 2. We expect that the code will allow us to accurately compute millions-of-atoms configurations while effectively carrying out computations in domains that have an aggregate of tens of thousands of atoms. This would put nanoscale system prediction and design within reach in the immediate future. The approach is developed in collaboration with scientists from the Materials Science Division at Argonne.

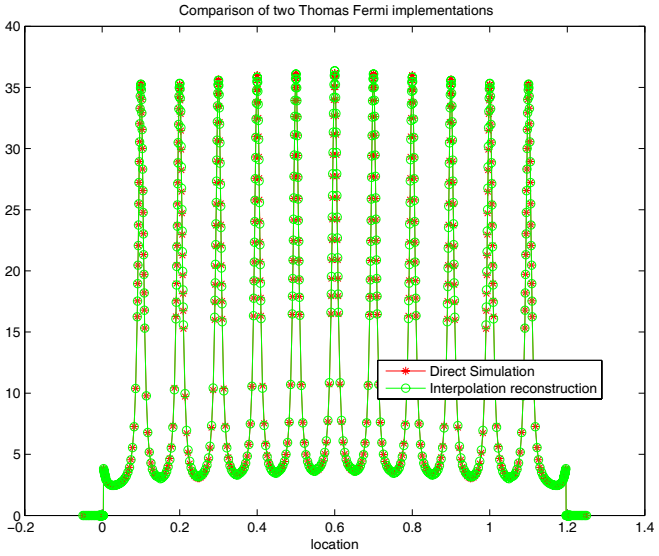

Figure 1: Comparison of the multiscale technique with the direct numerical simulation technique; electronic density plot

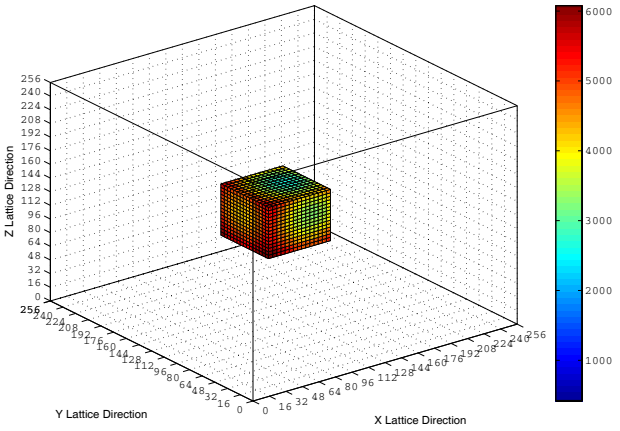

Figure 2: Computation of the forces on the nuclei with the MS-DFT code

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

Multiscale Design of Advanced Materials based on Hybrid *Ab Initio* and Quasicontinuum Methods

Michell Luskin, Richard James, and Ellad Tadmore, Univeristy of Minnesota
Michael Holst, Scott Baden, Bo Li, and John Weare, University of California, San Diego
Eric Bylaska [*] and Marat Valiev, Pacific Northwest National Laboratory

In this new program, starting in FY06, a multiscale method based on hybrid* ab initio *and quasicontinuum methods is being developed on a rigorous mathematical, physical, and chemical foundation. Rather than being based on classical interatomic potentials, this method will utilize quantum mechanics-based potentials capable of realistically describing the complex chemical bonding required to meet the design needs of advanced materials. Longer time scales will be achieved by a new accelerated* ab initio *molecular dynamics algorithm.

This project unites experts from mathematics, chemistry, computer science, and engineering for the development of new multiscale methods for the design of materials. To focus diverse approaches, we apply these methods to three families of materials that are of central interest to the mission of the Department of Energy (DOE): interfacial catalysts, metal organic framework (MOF) materials for hydrogen storage, and fusion materials.

Our approach is highly interdisciplinary, but it has two unifying themes: first, we utilize modern mathematical ideas about change-of-scale and state-of-the-art numerical analysis to develop computational methods and codes to solve real multiscale problems of DOE interest; and, second, we take very seriously the need for quantum mechanics-based atomistic forces, and base our methods on fast solvers of chemically accurate methods. We think that the extent of integration of chemistry and numerical analysis described here is unprecedented in the area of multiscale methods.

The most successful multiscale method in existence today, in terms of the actual problems it has solved and the supporting comparisons with fully atomistic computations and experiment, is the quasicontinuum (QC) method of Tadmor (Tadmor et al., 1996, www.qcmethod.com). The QC method has recently been integrated with density functional theory to allow chemically accurate predictions in the atomistic region. Yet QC methods, and most

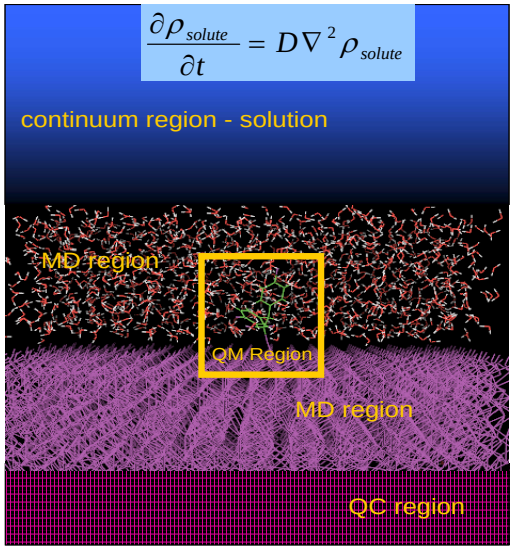

Figure 1. Depiction of a Hybrid Ab Initio *Quasicontinuum method.*

[*] 509-376-9114, eric.bylaska@pnl.gov

(if not all) multiscale methods in existence today, are largely *ad hoc*. We propose a novel numerical analysis of the QC method and use these results to develop a new multilevel, high-resolution version of QC with a strategy for using error estimation to drive mesh refinement and coarsening in QC.

Major extensions and capabilities are being added to QC, including the integration of QC with the more accurate quantum mechanics solvers of Weare, Valiev, and Bylaska; further development of "hot-QC;" and the extension of QC methods to complex lattices. We are developing new software techniques and tools for parallel implementation on advanced computer platforms to meet the challenges posed by the adaptive grids and multiphysics in QC and to enable the design of realistic advanced materials.

It is widely appreciated that to bring the Periodic Table to the design of materials one needs methods based as closely as possible on accurate quantum mechanics. We do this not by calling standard density functional theory (DFT) software, but by concurrent development of chemically accurate algorithms. The center piece of this effort is the *Ab Initio* Molecular Dynamics (AIMD) method. As with QC, we propose a pioneering mathematical analysis of the method with special attention to numerical analysis and efficient computing. With this approach, the needs and limitations of the multiscale methods can have a direct influence on the development of the next generation of chemically accurate first principles solvers.

In mathematical studies of multiscale methods, certain basic techniques underlie nearly all of the results. These studies most often begin by rescaling the larger and larger spatial (or time) scale problem onto a fixed domain, thereby giving a sequence of solutions that usually converge only weakly. The mathematics of dealing effectively with weak convergence is one of the most powerful and far-reaching themes in modern mathematics, and multiscale applications have motivated nearly all of this work. But most of these advances, such as methods of Gamma convergence or homogenization, have involved passing from one continuum theory to a larger scale one. Among the few studies of discrete theories, none to our knowledge involve chemically accurate methods at the atomic scale. This work brings these latest methods to the study of chemically accurate methods on large systems.

The organization of this research team maximizes the impact on the wider DOE community. Two of our participants, Bylaska and Valiev, work at Pacific Northwest National Laboratory (PNNL). As developers of chemically accurate first principles methods themselves, their own research is critical to the success of this project. They also bring a deep familiarity with the material systems we have identified, including both the behavior and severe requirements of their utilization. They will also serve as liaisons with the broader activities at DOE, including the rapid implementation of our ideas into the NWChem code and planned interactions with the new Institute for Interfacial Catalysis at PNNL. The latter gives us the opportunity of having direct experimental comparisons with our results.

For further information on this subject contact:
Dr. Gary Johnson, Program Manager
Mathematical, Information, and Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: 301-903-5800
gary.johnson@science.doe.gov

Northwest Consortium for Multiscale Mathematics and Applications – *"Educational Strategies and Critical Problems in the Thermo-mechanics of Materials"*

Rogene M. Eichler West [*], Kim Ferris, Benoit Mandelbrot,
Pacific Northwest National Laboratory
Alexander Panchenko, Sinisa Mesarovic, Hussein M. Zbib, and Ben Li,
Washington State University
R.E. Showalter, Oregon State University

Summary

The Northwest Consortium for Multiscale Mathematics is developing educational strategies to bring advances in multiscale mathematics into the graduate school classroom through the development of new curricula, textbooks, educational software, and summer programs. The motivation for these reforms comes from the need for researchers to be trained in emerging areas of multiscale mathematics so as to make significant contributions to critical problems in the thermo-mechanics of materials.

This three-year Northwest Consortium (NC) project began with planning activities in September 2005. We have made significant progress on the following four deliverables:

- A multiscale textbook
- A multifractals textbook
- A web-based educational software library
- A summer school in multiscale mathematics

The scope for the multiscale textbook is nearly complete. Topics were derived from four sources: 1) surveys of mathematical methods that have been developed and applied in association with principal investigators receiving federal funding from either the Department of Energy (DOE) or the National Science Foundation (a complete listing will publicly available online); 2) surveys of mathematical methods published since the inaugural issues of Multiscale Mathematics and Simulation (SIAM) and the International Journal for Multiscale Computational Engineering; 3) attendance at a series of multiscale workshops sponsored by the Institute for Pure and Applied Mathematics on the University of California, Los Angeles campus; and 4) individual communication with researchers. Work on individual chapters will begin in November in collaboration with the developers of the methods. The textbook is targeted for graduate level and is intended to gather and disseminate the state-of-the-art in multiscale approaches across all science domains.

The multifractals textbook will be authored by Benoit Mandelbrot, who is an internationally respected expert on multifractals. Plans to establish research and publications across multiple science domains in support of this textbook include: meetings between affiliated research groups and

[*] 509-375-4503, rogene.eichler.west@pnl.gov

Mandelbrot; training for affiliated researchers interested in multifractals via a semester-long seminar series (in conjunction with Washington State University, Tri-Cities); hands-on data workshops for affiliated researchers led by Mandelbrot; participation by affiliated researchers in the chapter review process and the alignment of the research projects as supporting evidence in the book.

While their research is not directly supported by this project, the educational opportunities to learn more about the application of multifractals will serve to unite affiliated researchers who have expressed their interest in applying multifractals to the following problem domains:

- Analysis of the spatial variability of soil properties, such as hydraulic conductivity and solute dispersivity
- Multiscale representation and analysis of radiation flux from clouds and ground objects in climate models
- Remote sensing of environmental pollutants, such as industrial stack plumes
- Identification of compact representations for proteomics signatures based on fractal analysis of mass spectrometry data
- Analysis and predictive correlations between energy price market data and energy production facilities
- Efficient simulation methods in molecular dynamics and the analysis of the resultant emergent structures in thin films.

This multiscale work on the topic of fractals continues to enhance outreach and pipeline activities for the recruitment of student interns at Pacific Northwest National Laboratory (PNNL). Fractals is an area that captures the imagination of young people and helps to encourage the consideration of careers in math and science. Two such interns joined PNNL last month and are learning JAVA programming and Web development to make contributions to the Web-based educational software library.

A preliminary version of the website is anticipated to be opened in the last quarter of CY2005. The structure is not yet fixed, and this early version will serve as the initial focal point for future discussions on our educational and summer school deliverables. The structure of the site is in place. We are gathering content from other researchers who are working on methods and applications in the multiscale mathematics field so that the website might become a virtual community spanning disparate fields of mathematical and scientific research. Mathematical software supporting the textbook and the summer school will be the focus of the next phase of development.

In preparation for the month-long summer school in 2007 and 2008, we will hold three week-long workshops during the 2005-2006 funding year. The workshop topics agreed upon were homogenization for discrete and continuum systems, computational issues in multiscale computing, and multifractals. Speakers have also been identified for all three workshops. An agreement has been made with Western Washington University to hold the multifractals workshop on their campus June 17–21, 2006.

Finally, members of the NC are developing timelines for the research and curriculum aspects of this project. These timelines will be presented during the Kick-off Meeting planned for the late October, when we will vet our approach and receive feedback from our DOE program manager.

For further information on this subject contact:
Dr. Gary Johnson, Program Manager
Mathematical, Information, and Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: 301-903-5800
gary.johnson@science.doe.gov

Mathematical Programming with Equilibrium Constraints – Expanding the Scope of Optimization

Mihai Anitescu* and Sven Leyffer* Argonne National Laboratory

Summary

What have electricity markets, toll-road pricing, and quasi-brittle fracture identification in common? They can all be modeled as optimization problems with equilibrium constraints. Researchers at Argonne National Laboratory have developed new theoretical insight and computational techniques to tackle this growing class of challenging problems.

Optimization, the art of finding the minimum value of an objective function subject to constraints, is one of the fundamental modeling paradigms of mathematics and computational science. It appears in virtually any area of science and engineering. Yet in spite of decades of advances in the theory and algorithms for optimization, important challenges remain, especially in the case where both the function and constraints are nonlinear, giving rise to the nonlinear optimization case. Overcoming these challenges is an important step toward creating robust software for many scientific computation applications.

Recently, an important class of problems that poses special challenges has been identified by the optimization community: mathematical programs with equilibrium constraints (MPECs). MPECs arise, for example, in the modeling of electricity markets as so-called Stackelberg games. These problems are characterized by a dominant producer (the leader) that can exercise marker power by anticipating the optimal responses of the smaller producers (the followers). The leader maximizes its profits subject to the optimal response of the followers. The result is an equilibrium constraint in the leader's optimization problem and, thus, an MPEC.

MPECs are computationally challenging problems because they violate standard stability assumptions. To make matters worse, early numerical experience was so disappointing that researchers concluded that classical nonlinear optimization algorithms would never solve MPECs reliably. This conclusion was unfortunate, since the set of applications described by MPECs is vast, including many diverse applications such as fracture identification, optimal control of cost-efficient robots, optimal design of contact configurations in mechanical and civil engineering, multiscale modeling in nanoscience, and data classification.

Recently, our group has shown that this gloomy prognosis was premature. By combining new theoretical insight and novel computational developments, we have been able to reliably and efficiently solve a large

* **Mathematics and Computer Science Division, (630) 252-4698, anitescu@mcs.anl.gov, leyffer@mcs.anl.gov**

class of MPECs. Our techniques routinely solve applications that are orders of magnitude larger and more complex than previously possible, extending the scope of this important new computational paradigm. An example of the solution of a medium-scale problem is presented in Figure 1, where the optimal contact region between an elastic membrane and a rigid obstacle, defined by an MPEC, is successfully solved by our methods.

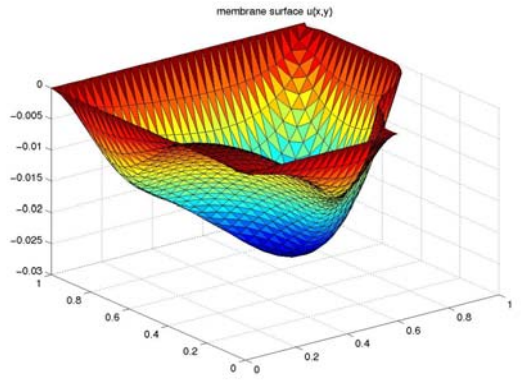

Figure 1: Minimal surface membrane over an obstacle

The equilibrium constraint models the contact area, where the membrane rests on the obstacle. This contact area typically has a complex geometric structure, shown in red in Figure 2.

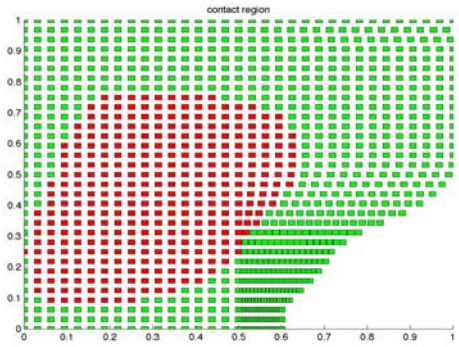

Figure 2: Obstacle contact area in red

Many popular nonlinear optimization codes have been modified to include our techniques and to expand their set of problems that can be reliably solved. Another indication of the increasing impact of our work is the fact that our five research papers have collected more than 100 citations (even though four have not yet appeared in the literature!).

In addition, we have collected over 150 MPECs in a test problem library, www.mcs.anl.gov/~leyffer/MacMPEC/, which has become the de facto benchmark for research on equilibrium constraints.

As a result of our work, we are now helping policy makers to better understand how electricity markets operate. Moreover, engineers are able to better predict how cracks in dams and other brittle structures develop.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

Interface Reconstruction in Multifluid Flow Simulations

Rao Garimella*, Vadim Dyadechko*, Blair Swartz*, and Mikhail Shashkov*
Los Alamos National Laboratory

Summary
We have developed an advanced method for reconstructing piecewise linear approximations of material interfaces in multimaterial flows. Results indicated that our method reconstructs smoother, more continuous and more accurate interfaces than other methods in its class.

Hydrodynamic simulation of multimaterial flows is important for understanding many problems such as droplet deposition, sandwich molding, underwater explosions, mold-filling in casting, and the behavior of microjetting devices.

An important feature of such flows is the interface between the materials, and it is often crucial to follow such interfaces during the simulation. Lagrangian simulations (where the mesh nodes move with the flow) automatically maintain interfaces but cannot handle large deformations and topology changes of these interfaces. On the other hand, Eulerian simulations (where the mesh nodes are stationary) must incorporate special procedures to keep track of the interfaces in the flow.

We have developed an advanced method for recovering a piecewise linear approximation of material interfaces in flow simulations given the volume fractions of materials in the cells of an unstructured mesh. The method incorporates several new techniques designed to make the reconstruction method more accurate, rapid, and robust. These include the careful selection and use of interface neighbor cells, and a topological consistency checking and repair algorithm for the interface that is designed to minimize fragmentation of the material regions being reconstructed. The volume of the materials is conserved exactly in the method.

The main steps of the procedure are:

Interface Estimation: A rough estimate of the interface is constructed using the volume fraction data specified on cells. The interface is represented by one line segment per cell.

Interface Smoothing: Interface segments are adjusted taking into account other interface segments in the neighborhood so that the resulting interface is as smooth as possible. Straight-line interfaces are typically recovered by this interface smoothing step.

Interface Topology Repair: The interface segments are adjusted so that they satisfy essential consistency requirements. The result of this is that the reconstructed material regions are continuous and do not have holes or fragments as far as possible.

---

* Mathematical Modeling and Analysis, (505) 665-2928, rao@lanl.gov, dyadechko@lanl.gov, bks@lanl.gov, shashkov@lanl.gov

(a)

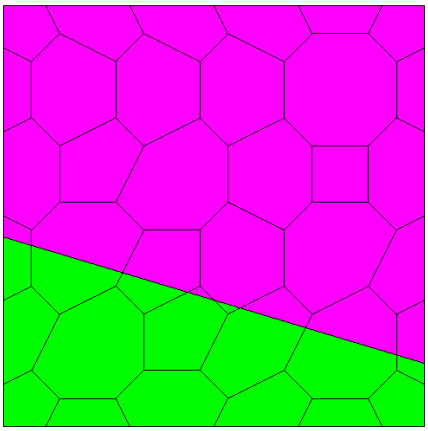

(b)

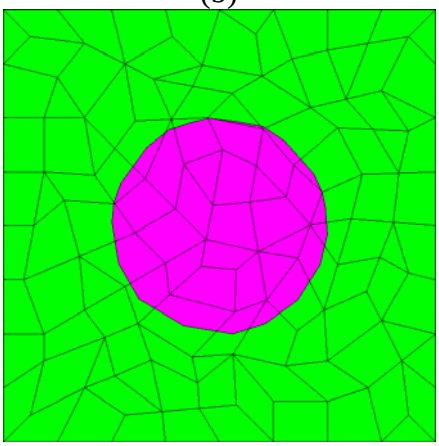

(c)

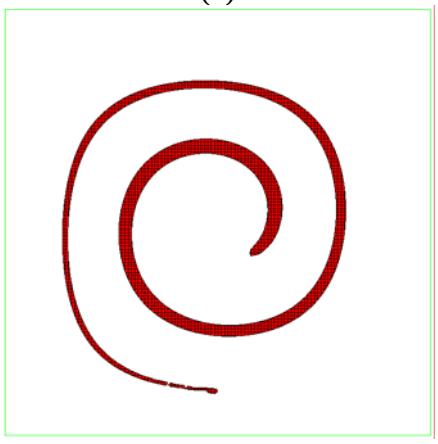

Interface reconstruction on structured and unstructured meshes (a) Straight Line Interface (b) Complex Interface (c) Vortex-in-box test of Rider and Kothe [RIDE98a]

Constrained Interface Smoothing: Alterations made to the interface in the repair step are smoothed, with the constraint that the topological consistency of vertices cannot be destroyed.

Interface Subdivision and Matching: The interface segment in each cell is subdivided into two. The two segments in the cell are then adjusted so that their slopes match the slopes of the appropriate interface segments in neighboring cells and the material volumes in the cell are conserved.

The procedure recovers complex interfaces more accurately than other interface reconstruction procedures, particularly exhibiting reduced fragmentation in the single vortex test.

References:

Rider, W. J. and D. B. Kothe, "Reconstructing volume tracking," *Journal of Computational Physics*, 141:112–152, 1998.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301) 903-5800
Gary.johnson@science.doe.gov

Bloch-Front Turbulence in a Periodically Forced Belousov-Zhabotinsky Reaction

Aric Hagberg*, Los Alamos National Laboratory

Summary

We have demonstrated, in a periodically forced oscillatory Belousov-Zhabotinsky reaction and in mathematical models, a mechanism for creating spatio-temporal disorder. The mechanism consists of the creation of spiral vortex pairs through a transverse instability of fronts in the vicinity of a nonequilibrium Ising-Bloch bifurcation. This is the first direct experimental evidence tying front instabilities to vortex nucleation and disorder.

Spatio-temporal disorder in extended systems commonly involves the spontaneous creation and annihilation of localized structures such as defects and vortices. The driving forces for the nucleation of defects and vortices are instabilities of periodic patterns or fronts. Defects in periodic patterns often result from the Benjamin-Feir-Newell instability, while spiral-vortex nucleation in bistable systems has been related to a front instability – the nonequilibrium Ising-Bloch bifurcation.

We have demonstrated, in a periodically forced oscillatory Belousov-Zhabotinsky reaction and in mathematical models, a mechanism for creating spatiotemporal disorder. The mechanism consists of the creation of spiral vortex pairs through a transverse instability of fronts in the vicinity of a nonequilibrium Ising-Bloch bifurcation. This is the first direct experimental evidence tying front instabilities to vortex nucleation and disorder.

We used an amplitude equation model, the forced complex Ginzburg-Landau equation,

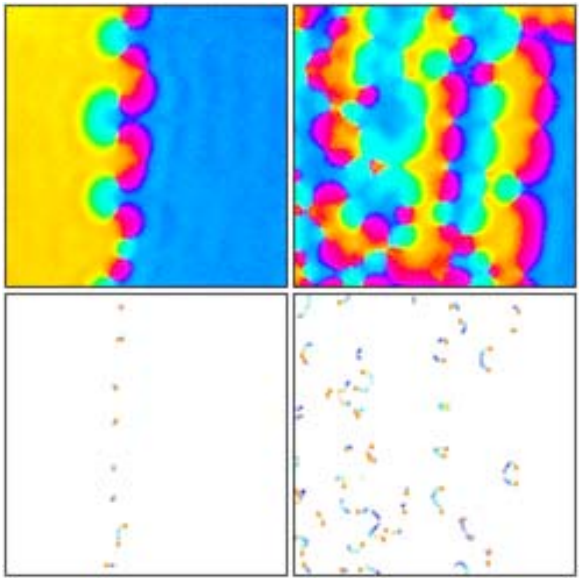

Figure 1: Experiments on a periodically forced Belousov-Zhabotinsky chemical reaction show front breakup into a state of spatiotemporal disorder involving continual events of spiral-vortex nucleation and destruction. The initial front of the oscillation phase is unstable to transverse perturbations and vortices form in pairs along the front line. The bottom frames show the position of the vortices, the core location of the spiral wave, as red dots with colored tails indicating the motion and direction of travel. Our amplitude equation models capture the spatiotemporal dynamics of the chemical system.

* Mathematical Modeling and Analysis Group, Theoretical Division, hagberg@lanl.gov, http://math.lanl.gov/~hagberg/

to reproduce the experimental observations with numerical solutions, and further described the mechanism for vortex creation with the normal form equations for a curved front line.

The normal form equations reduce the two-dimensional problem to the study of a one-dimensional set of equations for the curvature and velocity of the interface. With these equations we can capture the physics of a spiral-vortex nucleation event.

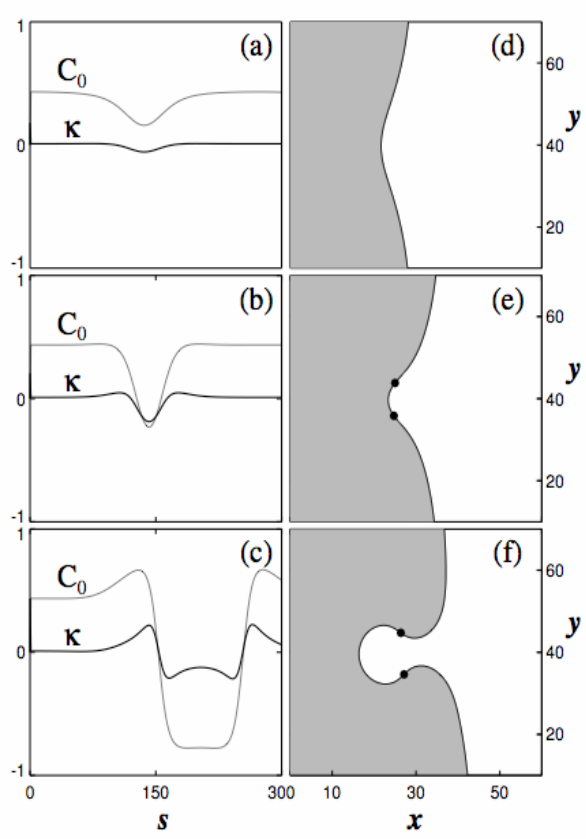

Figure 2: Nucleation of a single spiral-vortex pair in the front-line equations. Frames (a)-(c) show the front velocity κ and curvature C_0 vs the arclength s. Frames (d)-(f) are the corresponding representation in the laboratory coordinate frame. A small perturbation in the curvature grows and a portion of the domain reverses direction causing spiral-vortex pair to nucleate along the front line.

References:

Bradley Marts, Aric Hagberg, Ehud Meron, and Anna L. Lin, "Bloch-front turbulence in a periodically forced Belousov-Zhabotinsky reaction," *Phys. Rev. Lett.*, Vol. 93, p. 108305, 2004.

Co-investigators:

Bradley Marts, Anna L. Lin
Department of Physics
Duke University, Durham, NC 22708-0305

Ehud Meron
J. Blaustein Institute for Desert Research
Ben-Gurion University
Sede Boker Campus, Israel 84990

For further information on this subject contact:

Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301) 903-5800
Gary.johnson@science.doe.gov

Clean Air and Power Games

Sven Leyffer*, Argonne National Laboratory
Todd S. Munson*, Argonne National Laboratory

Summary

Can pollution control increase the market power of electricity producers? What are the interactions between pollution control and electricity markets? These questions are important for the design of an efficient electricity market. We investigate the ability of the largest producer in an electricity market to manipulate both the electricity and emission allowances markets to its advantage. We construct a computational model of the Pennsylvania - New Jersey - Maryland electricity market and show that the leader can gain substantial profits by withholding allowances and driving up emission allowance costs for rival producers.

Market power is defined as the ability of players in a market – producers and consumers, for example – to unilaterally or collectively maintain prices above the competitive level. The exercise of market power can result in price distortions, production inefficiencies, and a redistribution of income among consumers and producers. The electricity market is especially vulnerable to the exercise of market power by the producers for three reasons. First, short-term demands for electricity are very inelastic, largely because consumers are shielded from fluctuations in real-time prices. Second, network limitations lead to market separation if transmission lines are congested. Third, supply curves steepen when output is near capacity, implying that the marginal cost increases drastically in segments where the electricity price is determined during peak periods.

Pollution control regulation can significantly increase production costs in electricity markets. The NOx allowances program in the eastern United States, for example, is a cap-and-trade program administered by the U.S. Environmental Protection Agency (USEPA). The amount of NOx released into the atmosphere under this program is controlled by distributing allowances to the producers that must be redeemed to cover actual emissions. These allowances can be traded in a secondary market or banked for future use.

Market power can interfere with the efficiency of the emission control program, yielding higher costs for both emission control and commodity production. An example of such market power is the ability of producers to use allowances as a vehicle to affect the costs of rivals. The consequences of exercising market power can be complicated because of the interaction between the electricity and allowances markets. Empirical analysis of the 2000–2001 California power crisis, for example, suggests that in addition to demand growth, a shortage of hydropower, and excessive reliance on spot markets, some price increases were caused by a large producer that intentionally consumed more allowances than necessary, raising the costs

* Mathematics and Computer Science Division, (630) 252-4698, leyffer@mcs.anl.gov, tmunson@mcs.anl.gov

for rival producers that were short of allowances.

We construct a Stackelberg game for the Pennsylvania - New Jersey - Maryland Interconnection (PJM) electricity market. This model differs from other oligopolistic models in the following ways. First, interaction between the emissions and electricity markets is explicitly represented in the model. In particular, the allowances price is endogenously determined, as opposed to being an exogenous quantity as in other models. Second, the model is developed from the bottom up and is based on detailed engineering data for a power system with 14 nodes, 18 arcs, and 5 periods, shown in Figure 1.

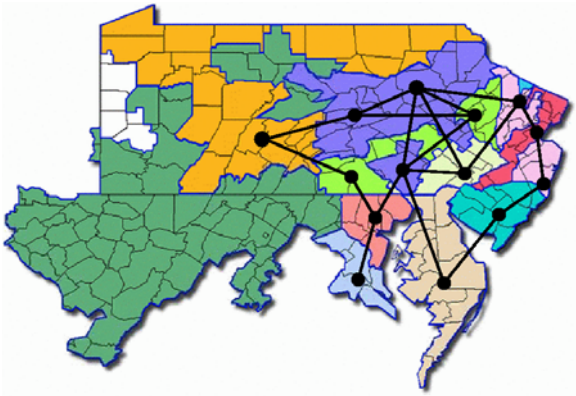

Figure 1: Network topology of PJM market

The data incorporated in the model includes heat rates, emission rates, fuel costs, location, and ownership for each generator. This approach allows for a more realistic estimation of the market power associated with the location of a generator in the network. Moreover, the power flow in the network is represented by a linearized direct-current load flow model in which the Kirchhoff current and voltage laws account for quadratic transmission losses.

The resulting model has 18,500 variables and 9,500 constraints, making it one of the largest Stackelberg games solved to date. A two-phase strategy was used to solve the problem that first finds a feasible iterate by solving a related complementarity problem and then applying an optimization algorithm that builds on recent developments in solving ill-conditioned nonlinear optimization problems.

The solutions to the Stackelberg game for the PJM electricity market show that the leader can gain substantial profits through the exercise of market power at the expense of other producers. Whether the withholding allowances strategy is profitable depends on (among other factors) the net position of the leader in the NOx allowances market. According to this model, some producers may be in a position to profit from withholding allowances; however, it is not optimal for other producers to undertake such practices.

This computational experience is promising for policy modelers interested in investigating the complicated interactions among imperfectly competitive markets.

For further information on this subject contact:
Dr. Gary M. Johnson, Program Manager
Mathematical, Information & Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301)903-5800
Gary.johnson@science.doe.gov

Exploratory Data-Intensive Computing for Complex Biological Systems

T.P.Straatsma [*], Pacific Northwest National Laboratory
N.F.Samatova [#], Oak Ridge National Laboratory

Summary

The goal of this "BioPilot" project is to explore new mathematics and computational methodologies in support of biological computing problems that rely on especially large and complex sets of data. This involves the development of new algorithms and prototype codes that take optimal advantage of new, globally addressable memory architectures. This exploratory project includes performance evaluations for various types of problems, and a determination as to the feasibility of these architectures for the large-scale computational needs of the biology community.

The revolution in the computer science and information technology world is happening at the same time as the revolution in the biological sciences. Many types of biological computing problems rely upon very large sets of data that must be shared effectively across many processors. Globally addressable memory architectures provide an opportunity to solve new large-scale biological computing problems that do not seem suitable for other types of architectures.

Three biology focus areas where advanced data intensive computing capabilities have been explored are: (1) the analysis of mass spectrometry data using production-scale proteomics measurements, (2) modeling of protein structure, molecular docking, and the dynamics of molecular machines, and (3) the development of algorithms and prototype codes for reconstruction of network topology and stochastic network simulation.

This project involves the development of mathematics and algorithms and prototype codes for these environments, performance evaluations for various types of problems, and a determination as to the feasibility of these architectures for the large-scale production needs of the biology community.

Peptide identification

The false-positive peptide identification rate was decreased by an order of magnitude by enhancing rigorous statistical analysis methods for high-throughput mass spectrometry. Dramatic gains in *de novo* sequencing performance were achieved due to advanced probabilistic scoring functions. The robustness of identification methods for medium- to low-quality MS/MS spectra was increased by developing novel sequence optimization methods. Parallel, scalable to thousands of processors, computations of scoring functions for peptide identification was prototyped on globally addressable memory architectures. Novel underlying memory layouts were designed to efficiently extract complex relationships between the peaks in the spectra for improving correct peptide identification rate. Utilization of memory-intensive mass indices accelerated

[*] 509-375-2802, tps@pnl.gov
[#] 865-241-4351, samatovan@ornl.gov

the database search step 100-200 times. Time-to-solution for the entire proteome comparison against the nonredundant protein database was reduced from weeks to minutes by utilizing a novel shared memory programming model and by hiding the scalability gap between computation and I/O.

Biomolecular and biophysical modeling

A new molecular trajectory analysis implementation has been developed that is based on keeping complete trajectories in core, leading to a significant increase in efficiency and scalability both with trajectory size and number of processors used in the parallel execution.

The novel methodology was developed to speed up the energy calculations by 20-50 times for the large range biophysical modeling problems. Pre-calculated Van der Waals interactions for a large set of the specific spatial arrangements between various small fragments of biomolecular systems are stored in large shared memory data structures.

Biological network simulation

A parallel exact stochastic algorithm was developed with the aim of speeding the stochastic simulations so that large sized biological networks can be investigated. Reduced inter-process communication latency was demonstrated to be critical for the efficient implementation of these communication-intensive applications. With the growing processor speed, the need for minimizing the latency will be increasing.

We have also developed scalable mathematical models that can be utilized in testing the parallel stochastic algorithms.

Network inference and analysis

New highly parallel implementations of the BLAST sequence alignment tool (ScalaBLAST and PioBLAST) have been developed. Novel genomic context analysis methods were prototyped, including the implementation of a parallel complete-linkage hierarchical clustering algorithm on a globally addressable architecture. Memory-intensive clique-based analysis of gene expression networks using large shared memory machines enabled to expand vastly the group of known pathway genes responsible for genetic variation in behavioral phenotypes of genomes under study.

Impact and direction

By exploring the feasibility of globally addressable memory architectures for the large-scale production needs of the biology community, we not only have learned the important aspects of the unique features of data-intensive problems in systems biology from the algorithmic and architectural point of view, but also have provided proof-of-principle demonstrations that are likely to lead to robust solutions.

The work in the BioPilot has a broad impact on many data-intensive applications in systems biology important to the DOE mission. More importantly, the technologies being developed are of a more general nature that can be applied beyond the target applications chosen for the first year proof-of-principle demonstration. The focus of the project is to continue to generalize and to make robust the key data-intensive biological applications with the pressing needs for computational advancements, and organize them into a user-friendly framework widely accessible to the broad scientific community.

For further information on this subject contact:
Dr. Gary Johnson, Program Manager
Mathematical, Information, and Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: 301-903-7774
gary.johnson@science.doe.gov